Project Status Report

High End Computing Capability
Strategic Capabilities Assets Program

7 March 2012

Dr. Rupak Biswas – Project Manager
NASA Advanced Supercomputing (NAS) Division
NASA Ames Research Center, Moffett Field, CA
Rupak.Biswas@nasa.gov
(650) 604-4411

Significant Accomplishments – Current and Future

Current

- PBSPro Update on Pleiades Results in Additional Available Computational Cycles
- HECC Develops New Tools to Enhance Pleiades Reliability
- New Method Detects Use of Uninitialized Floating Point Variable in Application
- HECC Recovers Critical Air Traffic Management Data
- Proactive Monitoring of User Jobs Improves System Utilization for Research Scientist
- User Services Team Completes RSA SecurID Token Replacement
- Webinars Provide Training for Both HECC Users and Staff
- HECC Supports ECCO3 for Studies of Ocean-Ice Interactions in Earth System Models
- Modeling & Simulation Support for NASA's Space Launch System
- NAS Managers Showcase High-End Computing Capabilities at NITRD Symposium

Future

- Pleiades SandyBridge Expansion

PBSPro Update on Pleiades Results in Additional Available Computational Cycles

- The HECC Supercomputing Systems team has deployed PBSPro 11.3, which provides feature enhancements and improvements to the batch scheduler on Pleiades.
- After the upgrade, the scheduling cycle time improved by a factor of 2x, resulting in faster start times for batch jobs and a much more responsive experience for users requesting interactive jobs.
- A significant benefit with PBSPro 11.3 is the addition of the "shrink-to-fit" feature, which enables jobs to start if they can complete useful computations before the start of a system dedicated time period.
- Shrink-to-fit allows HECC to recover nearly 600,000 System Billing Units (SBUs) for each dedicated time; previously, these SBUs would be unused as the system was drained for dedicated time.

Mission Impact: Enhancements to the PBSPro batch scheduling system on Pleiades enable more computational cycles to be available for NASA science and engineering users.

Figure: Since the PBSPro upgrade on the Pleiades supercomputer, the scheduling time for users' batch jobs has improved by 2x.

POC: *Bob Ciotti, bob.ciotti@nasa.gov, (650) 604-4408, NASA Advanced Supercomputing Division*
Davin Chan, davin.s.chan@nasa.gov, (650) 604-4613, NASA Advanced Supercomputing Division, Computer Sciences Corp

HECC Develops New Tools to Enhance Pleiades Reliability

- The HECC Supercomputing Systems team recently developed three new tools to improve reliability and problem detection on Pleiades.
- The first tool is run on all nodes prior to a job execution. It detects insufficient free memory on any given compute node, automatically corrects the issue, and notifies HECC staff; this improves the user experience by ensuring that enough memory is available for batch jobs before allocating the node.
- The second tool is run automatically by the batch scheduler on completion of batch jobs to disable nodes that are left in a state with insufficient system disk space; this prevents a node from continuously accepting and then immediately rejecting new batch jobs due to insufficient space.
- A third tool automatically generates metrics from system logs in real time, which are then graphically displayed to assist staff in detecting trends and anomalies on the system.

Mission Impact: Development of new HECC tools improves the reliability and problem detection capability of Agency supercomputing systems, enabling these resources to be used more efficiently.

Figure: New tools improve the operation of HECC resources such as the Pleiades supercomputer, and enable users to be more productive.

POC: *Bob Ciotti, bob.ciotti@nasa.gov, (650) 604-4408, NASA Advanced Supercomputing Division*
Davin Chan, davin.s.chan@nasa.gov, (650) 604-4613, NASA Advanced Supercomputing Division, Computer Sciences Corp

New Method Detects Use of Uninitialized Floating Point Variables in Application Codes

- The HECC Application Performance & Productivity (APP) team developed a methodology using Intel compiler options and library-based interception of memory allocation to reliably detect the use of uninitialized floating point variables.
- Uninitialized variables result in garbage data being used in the computations, calling into question the correctness and reliability of the computer simulations.
- The APP team determined that the commonly used Intel compiler options meant to detect such uninitialized variable errors, failed to do so for a large class of typical user code constructs.
- Using the HECC-developed method, floating point exceptions are forced when uninitialized floating point data is accessed; the Intel debugger can then be used to pinpoint the line of code where this occurred, allowing the developer to locate the source of the error.
- Two production codes have already successfully made use of the new methodology to detect occurrences of uninitialized floating point variables.
- Code developers can incorporate the method into their test and validation suites to prevent such errors thus increasing the reliability of their codes.

Mission Impact: Facilitating detection of programming errors in complex production codes increases reliability of simulation results across all NASA Mission Directorates.

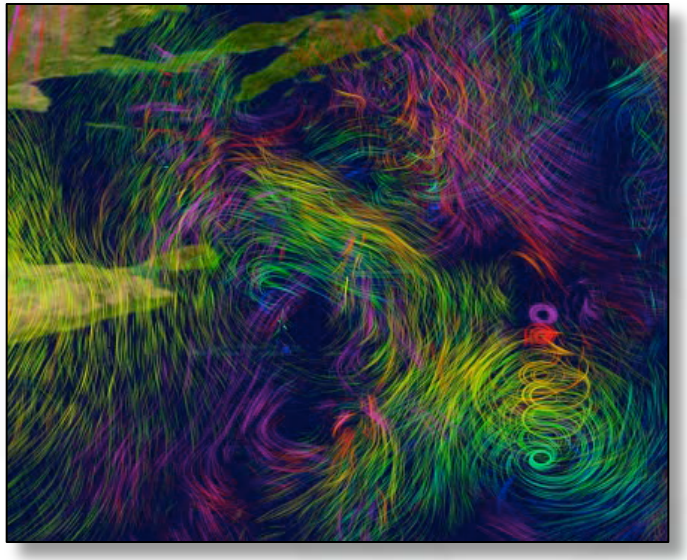

Figure: Finding uninitialized variables by hand in large production codes, such as in the example Earth science application shown above, is extremely difficult. Using such undefined variables results in behavior that can change when the execution environment is perturbed in a seemingly harmless way.

POC: *David Barker, david.p.barker@nasa.gov (650) 604-4292, NASA Advanced Supercomputing Division, Computer Sciences Corp.*

HECC Recovers Critical Air Traffic Management Data

- HECC provides a Disaster Recovery (DR) data storage capability for several important projects, including NASA's Next-Generation Air Traffic Management (ATM) research project and Ames Multi-Mission Operations Center.
- Recently, the ATM lost data in the Dynamic Weather Routes project on their internal systems; HECC systems experts restored the data from DR backups and enabled ATM to resume normal operations in a short period of time.
- HECC stores two duplicate copies of the ATM Disaster Recovery data in two separate computing facilities that are located 2 km apart; this ensures a high probability of data availability of HECC archive data.
- This backup storage capability averted a scenario that would have adversely affected the ATM project.

Mission Impact: The Disaster Recovery capability allows important NASA projects to leverage HECC Project resources and expertise in high-performance computing and data storage.

Figure: *To help air traffic controllers keep up with the anticipated heavy workload, NASA is developing several advanced automation tools that will provide controllers with more accurate predictions about the nation's traffic flow, weather, and routing.*

POC: *Bob Ciotti, bob.ciotti@nasa.gov, (650) 604-4408, NASA Advanced Supercomputing Division*
Davin Chan, davin.s.chan@nasa.gov, (650) 604-4613, NASA Advanced Supercomputing Division, Computer Sciences Corp

Proactive Monitoring of User Jobs Improves System Utilization for Research Scientist

- Recently, the APP team helped improve system utilization for HECC user Marissa Vogt, Dept. of Earth and Space Sciences, UCLA, who had submitted over 100 jobs, each requesting 16 nodes for 3 to 4 hours to execute a set of 128 one-core runs.
- Because of the huge backlog of queued jobs waiting for Pleiades resources, Vogt's jobs were languishing in the wait queue for over 15 hours, and could not be run as they would have interfered with the start of a higher-priority job.
- HECC's active monitoring of Vogt's past resource utilization showed that each job was using only 3 out of the 16 nodes and took only 24 minutes to run (only 21 of 128 one-runs were being executed).
- The improvements helped both the efficiency of Vogt's PBS jobs and their throughput by packing more serial jobs together to use all the cores and by reducing the requested wall-time to 45 minutes.

Mission Impact: HECC's active monitoring for efficient utilization of Pleiades resources is important to obtain optimal throughput for users' jobs. This is especially important, with a large payoff, at times when a huge backlog of jobs is waiting for resources, as has been the case for the past 2 months.

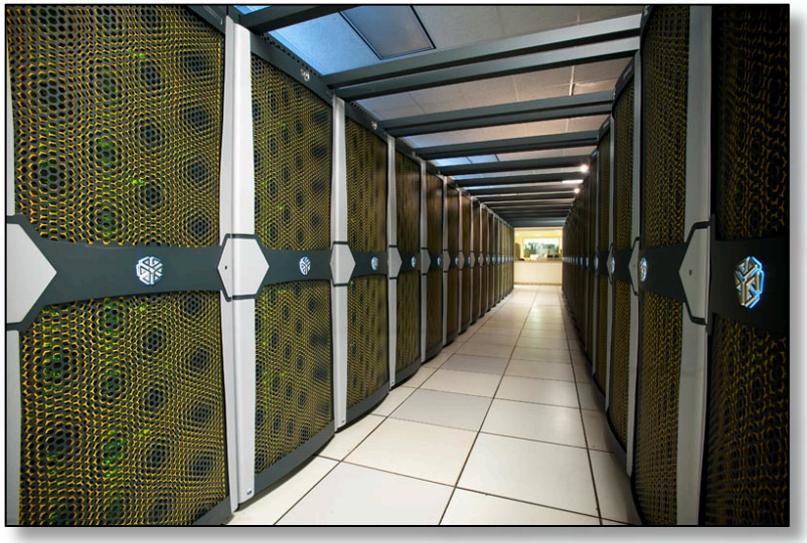

Figure: The HECC project will augment the computational capacity of the Pleiades supercomputer in April 2012. The addition of 24 SGI nodes (comprising Intel® Xeon® processors with 16 cores and 32 GB of memory per node) will help further alleviate the current large backlog of jobs.

POC: *Johnny Chang, johnny.chang@nasa.gov, (650) 604-4356, NASA Advanced Supercomputing Division, Computer Sciences Corp.*

User Services Team Completes RSA SecurID Token Replacement

- The HECC User Services team completed replacement of approximately 800 tokens, part of an Agency-wide project required due to a security breach at RSA SecurID.
- During this process, all active NASA Advanced Supercomputing (NAS) accounts were reviewed for accuracy and revalidated.
- Once the analysis was complete, Control Room personnel managed the distribution process over a period of two weeks.
- Tokens were mailed out in batches of 200 so that the Control Room could effectively manage the number of phone calls to activate user tokens.
- During this process, over 250 user accounts were deactivated and disabled for projects that had been completed.

Mission Impact: RSA SecurID tokens are at the foundation of authenticating to the HECC environment, providing continuous and secure access to resources and data.

Figure: Sample of the approximately 800 RSA SecurID Tokens replaced for NAS users

POC: *Leigh Ann Tanner, leighann.tanner@nasa.gov, (650) 604-4468, NASA Advanced Supercomputing Division, Computer Sciences Corp.*

Webinars Provide Training for Both HECC Users and Staff

- In November 2011, HECC initiated a series of monthly, web-based training seminars for the user community; to date, 4 different sessions have been presented, all focusing on effective use of HECC resources.
- The number of participating users has already increased to the point that, in February, HECC upgraded its WebEx service to accommodate up to 100 participants at a time (up from 25); numerous users have expressed appreciation for the new series.
- On February 21, HECC also conducted staff training on "System Monitoring using NAGIOS and the HUD," using the same webinar process; remote participants from the National Oceanic and Atmospheric Administration, and Alabama Supercomputing Center also attended.
- As with all other webinars, the latest session was recorded for use in training new HECC staff. HECC is also making this recording available to NCCS staff.

Mission Impact: Providing live and archived web-based training on best practices for utilizing NASA supercomputing resources allows users from all Mission Directorates to make more effective use of the systems—having a direct impact on their project milestones, at a very low cost.

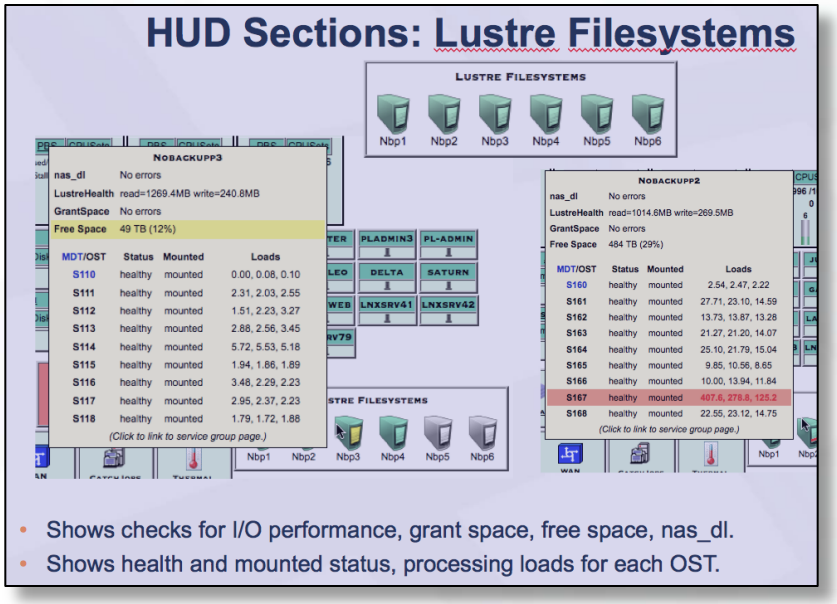

Figure: Slide from a staff training seminar on the Heads Up Display (HUD), delivered on February 21st. About a dozen people attended the training session, including five who participated remotely via WebEx.

POC: *Robert Hood, robert.hood@nasa.gov (650) 604-0740*
NASA Advanced Supercomputing Division, Computer Sciences. Corp.

HECC Supports ECCO3 for Studies of Ocean-Ice Interactions in Earth System Models

- HECC resources support a broad spectrum of Earth Science research projects, for example, Estimating the Circulation and Climate of the Ocean (ECCO) series of projects.
- The ECCO, Phase III (ECCO3) project aims to improve the representation of ocean-ice interactions in Earth System models.
- The specific objectives of ECCO3 are to study:
  1. Origin and evolution of water masses near polar ice sheets from high-resolution state estimates, numerical simulations, and adjoint sensitivity computations;
  2. Scientific basis for decadal climate predictability;
  3. Reduction of uncertainties in sea level projections through improved modeling of ice sheets and ocean-ice interactions.
- The nonlinear optimizations solved by ECCO3 are among the largest ever undertaken, involving billions of observations and control parameters and trillions of predictive model variables.
- HECC supercomputing resources allow the ECCO3 project team to run the huge minimization problems that must be solved in order to constrain the numerical model with observations.

Mission Impact: The availability of HECC supercomputing resources enables ECCO3 scientists to generate solutions that are improving estimates and models of the ocean carbon cycle in support of the Science Mission Directorate.

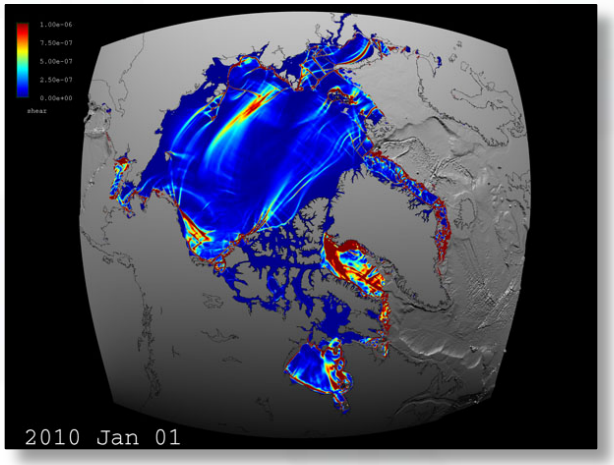

***Figure**:* Sea-ice shear (s-1) on January 1, 2010 in a coupled ocean and sea-ice simulation carried out by the ECCO3 project. The simulation has horizontal grid spacing of 4 kilometers, and is being used for comparisons with satellite retrievals from the RADARSAT Geophysical Processor System. (Gunnar Spreen, NASA/JPL; Tim Sandstrom, NASA/Ames)

***POC:** Dimitris Menemenlis, menemenlis@jpl.nasa.gov, (818) 354-1656, Jet Propulsion Laboratory, California Institute of Technology*

Modeling & Simulation Support for NASA's Space Launch System

- HECC resources and services enable fast and efficient turnaround time for modeling and simulation experts who perform important computational fluid dynamics (CFD) simulations of the Space Launch System (SLS) at key points during the ascent trajectory.
- These simulations predict aerodynamic performance, loads, and pressure signatures for different design variations, and include:
  - Initial shape trade studies to help assess and compare alternate designs developed at several NASA centers;
  - Inviscid aerodynamic performance characterization for both crew and cargo versions of SLS vehicle designs;
  - Viscous analyses of an early SLS design concept;
  - Computation of line loads and surface pressure signatures throughout ascent for preliminary designs.
- The Pleiades supercomputer enables extensive aerodynamic databases, comprised of hundreds to thousands of simulations, to complete in < 1 week.
- Databases of >1,000 inviscid simulations (slightly less computationally intensive than viscous solutions) complete in ~3 days using 12 cores per simulation.

Mission Impact: The availability of HECC supercomputing resources enables engineers supporting SLS to rapidly examine multiple options quickly exploring design spaces.

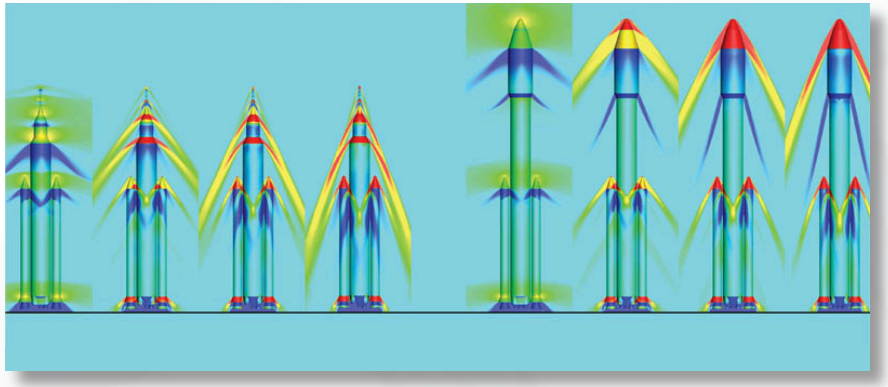

Figure: Contour plots of pressures on the vehicle surface and on a cutting plane for two early candidate crew and cargo launch vehicle designs at four Mach numbers. Michael Barad, NASA/Ames

POCs: *Cetin Kiris, cetin.c.kiris@nasa.gov, (650) 604-4485, Jeffrey Housman, jeffrey.a.housman@nasa.gov, (650) 604-5455, NASA Ames Research Center*

NAS Managers Showcase High-End Computing Capabilities at NITRD Symposium

- NAS Division managers joined other Agency leaders in advanced information technology (IT) to participate in the 20th anniversary symposium for the interagency Networking and IT R&D (NITRD) Program on Feb. 15–16, 2012, in Washington, D.C.
- NASA's exhibit featured an engaging poster, created by HECC staff, highlighting mission impacts from the Agency's investments over the past 20 years in advanced IT, including high-end computing, large-scale networking, and software development.
- The exhibit also featured handouts and videos summarizing NASA's aerospace and science mission activities enabled by IT R&D activities and capabilities.
- The NITRD Program is the nation's primary source of federally funded, revolutionary breakthroughs in advanced information technologies such as computing, networking, and software.

Mission Impact: Participating in interagency IT R&D coordination enables NASA to benefit from innovations by the federal community, and to share NASA innovations more broadly.

Figure: NASA poster showing specific examples of the impacts of Agency investments in advanced IT: Space Shuttle, future spaceflight, robotic exploration, aeronautics research, Earth science, and space science. The Agency's NITRD exhibit also featured videos of NASA's aerospace and science mission activities that have been enabled by IT R&D.

POC: *Bryan Biegel, bryan.biegel@nasa.gov, (650) 604-0171, NASA Advanced Supercomputing Division*

Status of Requests for NAS Computer Accounts by non-U.S. Citizens

- Requests approved: 13; New requests received: 6; Requests waiting: 1.
- One request was rejected because the applicant is an Iranian citizen. This is the first time a request has been rejected. (One other rejection was reversed upon appeal.)
- Average wait times have continued to decrease.
- The one user who is waiting has been waiting less than one month.

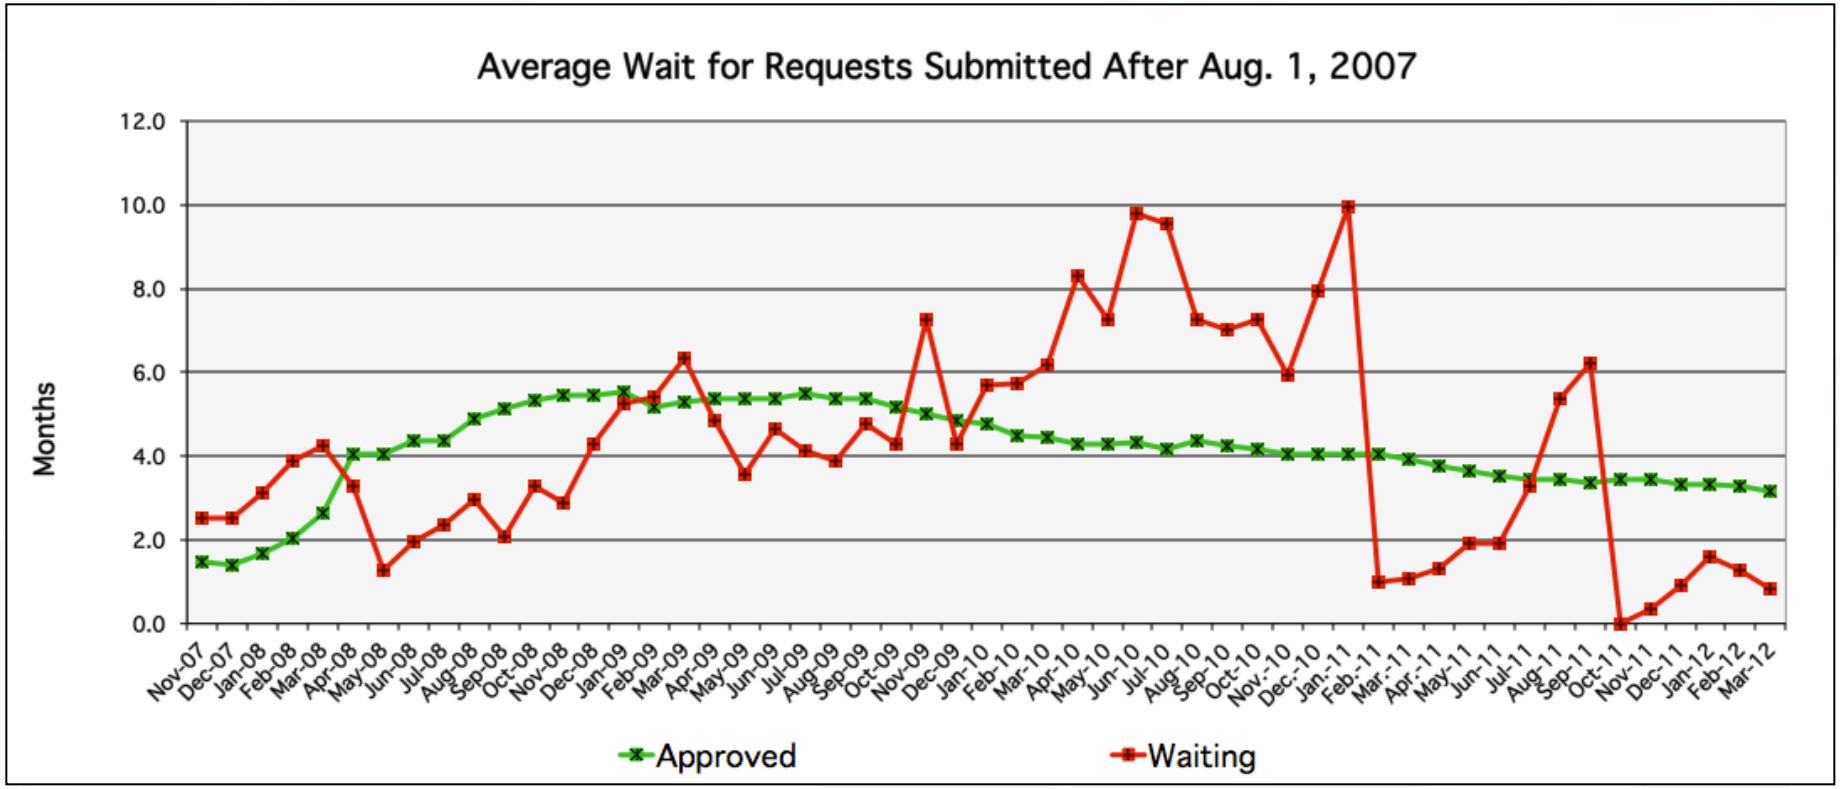

HECC Facility Hosts Several Visitors and Tours in February 2012

- HECC hosted 8 scheduled tour groups in February; guests received an overview of the HECC Project, demonstrations of the hyperwall visualization system, and tours of the computer room floor. Guests this month included:
  - NASA Chief of Staff David Radzanowski, along with Associate Director of Kennedy Space Flight Center Kelvin Manning, was briefed on the HECC Project and NASA Advanced Supercomputing (NAS) facility by Rupak Biswas, Bryan Biegel, Cetin Kiris, and Bill Thigpen.
  - Participants of NASA's Mid-Level Leaders Program, who were hosted by Ames during the 4th module of the program, learned about the Agency-wide missions being supported by Pleiades, and viewed scientific results on the hyperwall-2 visualization system.

Figure: NAS Division Chief Rupak Biswas (standing at rear) presents an overview of science and engineering projects run on HECC resources. From left: NAS Deputy Division Chief Bryan Biegel, Associate Director of Kennedy Space Center Kelvin Manning, Ames Associate Director Steve Zornetzer, NASA Chief of Staff David Radzanowski, and Ames Exploration Technology Directorate Chief Eugene Tu (in foreground).

POC: *Gina Morello, gina.f.morello@nasa.gov, (650) 604-4462, NASA Advanced Supercomputing Division*

Presentations and Papers

- "Modeling the Mesoscale Eddy Field in the Gulf of Alaska," Peng Xiu et al, *Deep Sea Research Part I: Oceanographic Research Papers,* Elsevier, available online 1 February, 2012.*
  http://www.umeoce.maine.edu/xiupeng/docs/Xiu_DSR1_2012.pdf
- "Two Earth-sized Planets Orbiting Kepler 20," F. Fressin, G. Torres, J.F. Rowe, D. Charbonneau, C.E. Henze, et al., *Nature, Volume 482, Pages:195–198, 09 February 2012.*
  http://www.nature.com/nature/journal/v482/n7384/full/nature10780.html

** HECC provided supercomputing resources and services in support of this work*

News and Events

- **NASA Scales Space Weather Simulation to 25K Cores on SGI Pleiades InfiniBand Cluster**, *SGI press release*, Feb. 27, 2012 – Highlights work by HECC user Homa Karimabadi (UC San Diego), who was able to utilize 25,000 cores on Pleiades to run a space weather simulation. Picked up many media sources, including HPCwire, InsideHPC, and Scientific Computing

  http://www.sgi.com/company_info/newsroom/press_releases/2012/february/nasa.html

HECC Utilization

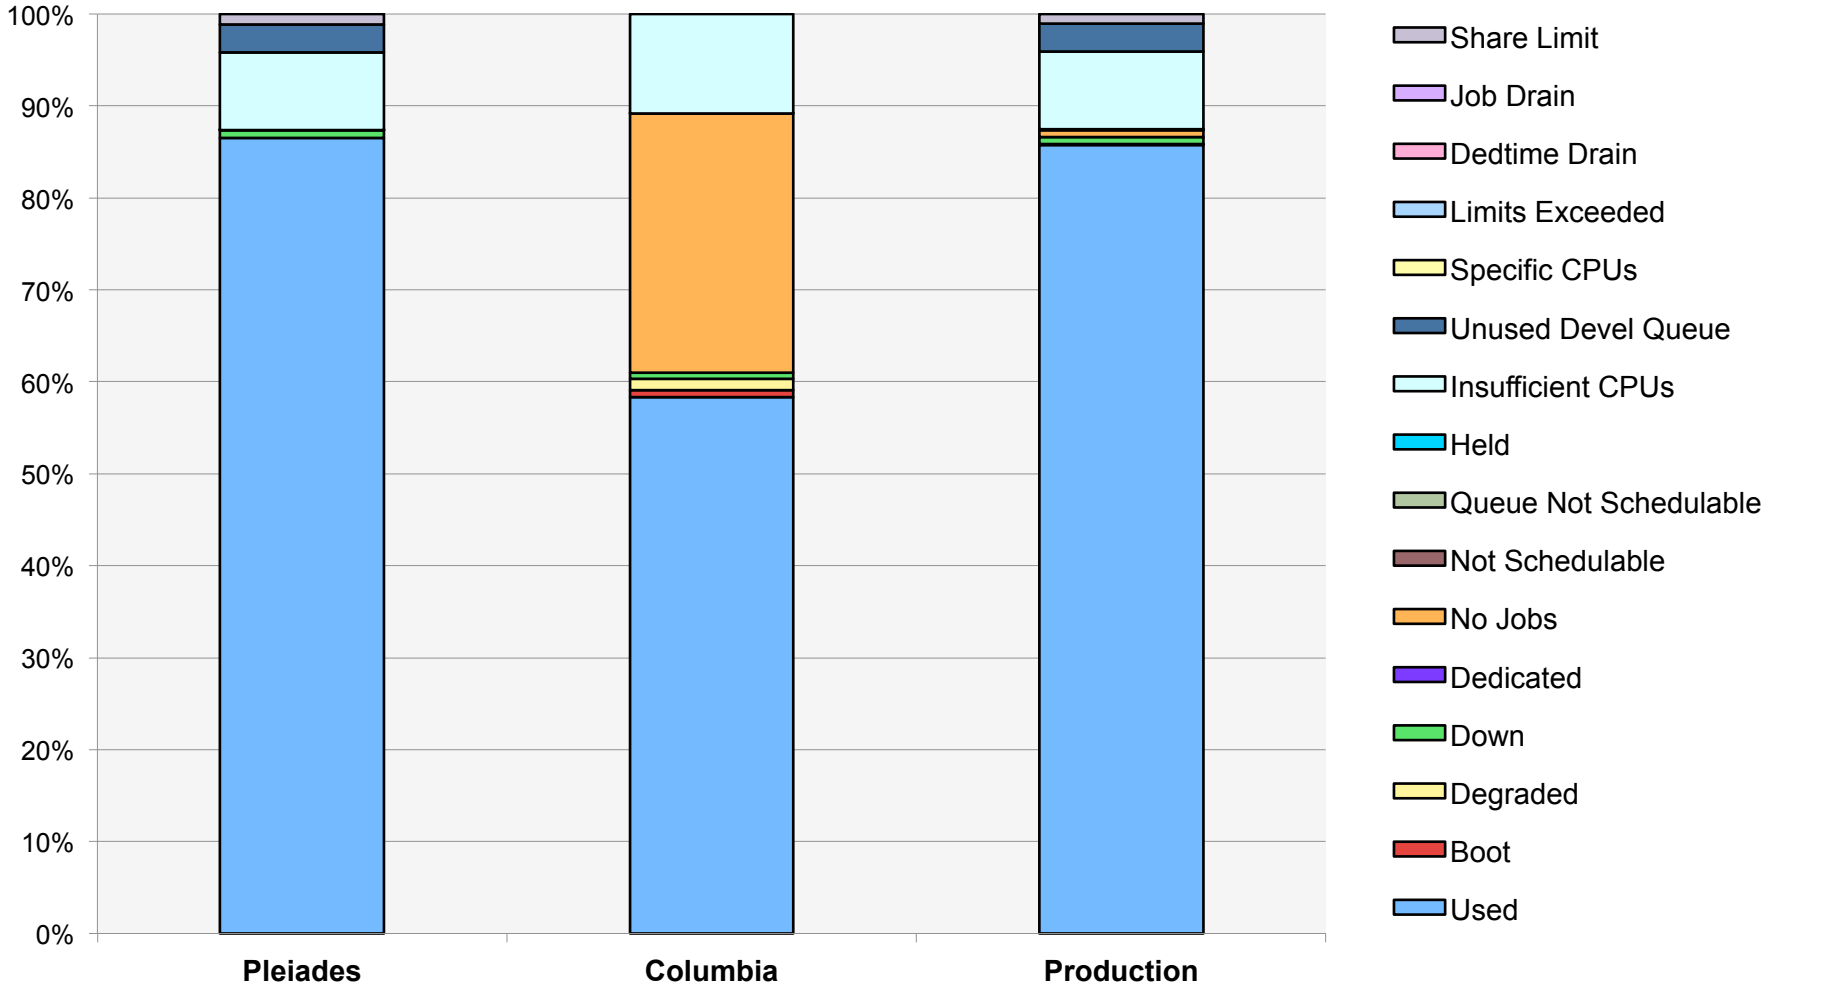

HECC Utilization Normalized to 30-Day Month

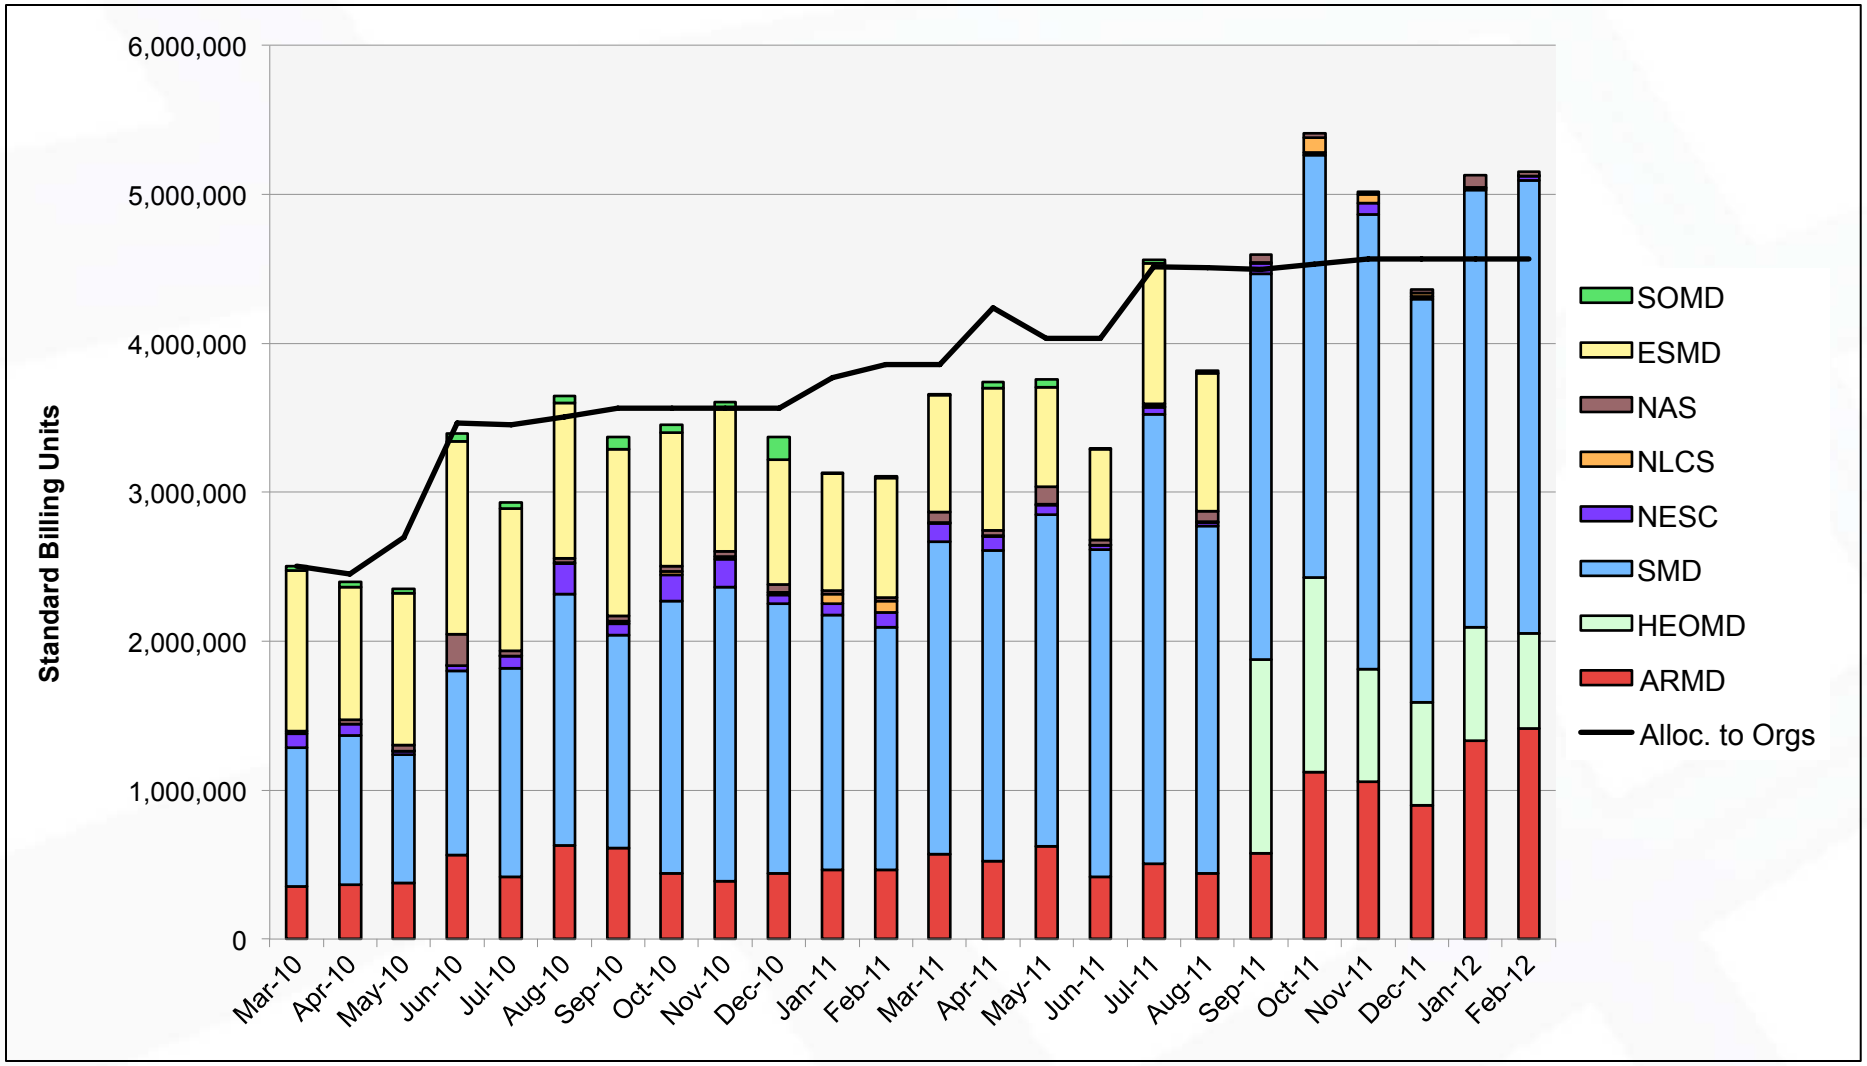

HECC Utilization Normalized to 30-Day Month

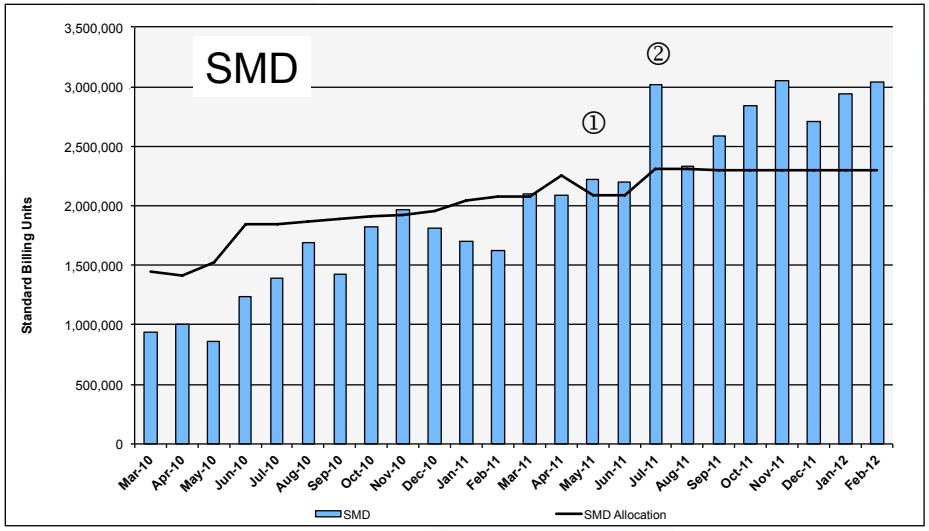

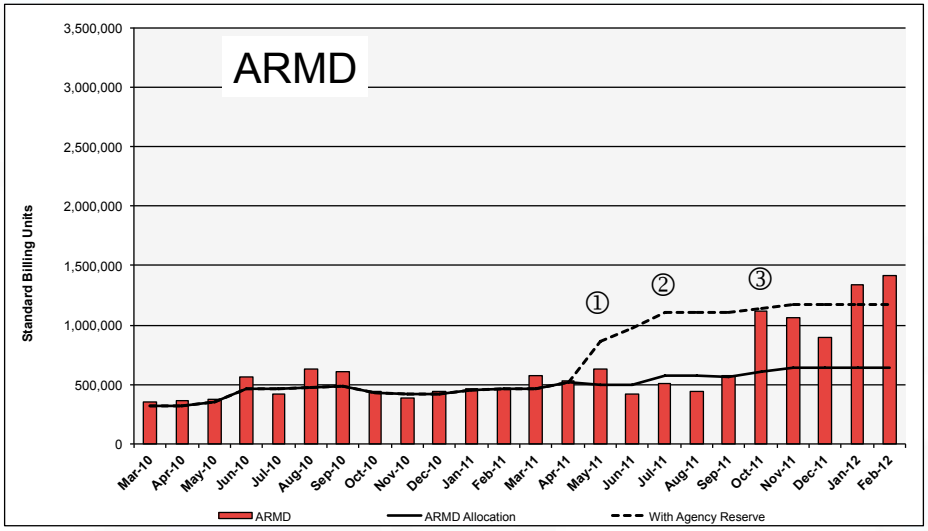

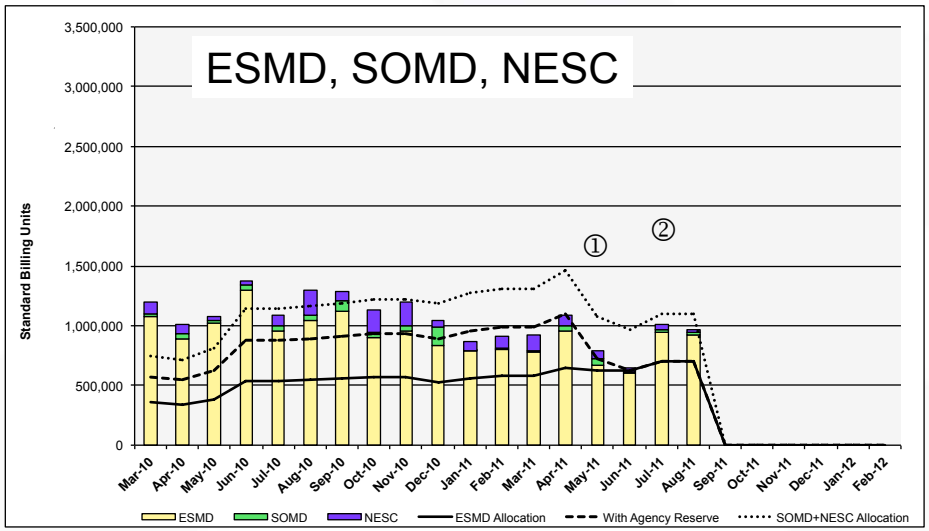

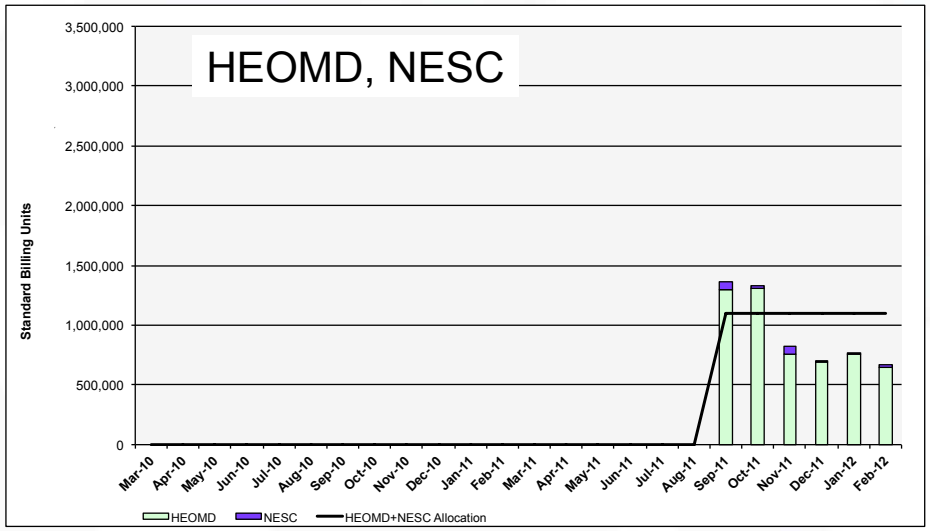

① Allocation to orgs. decreased to 75%, Agency reserve shifted to ARMD ② 14 Westmere racks added ③ 2 ARMD Westmere racks added

Tape Archive Status

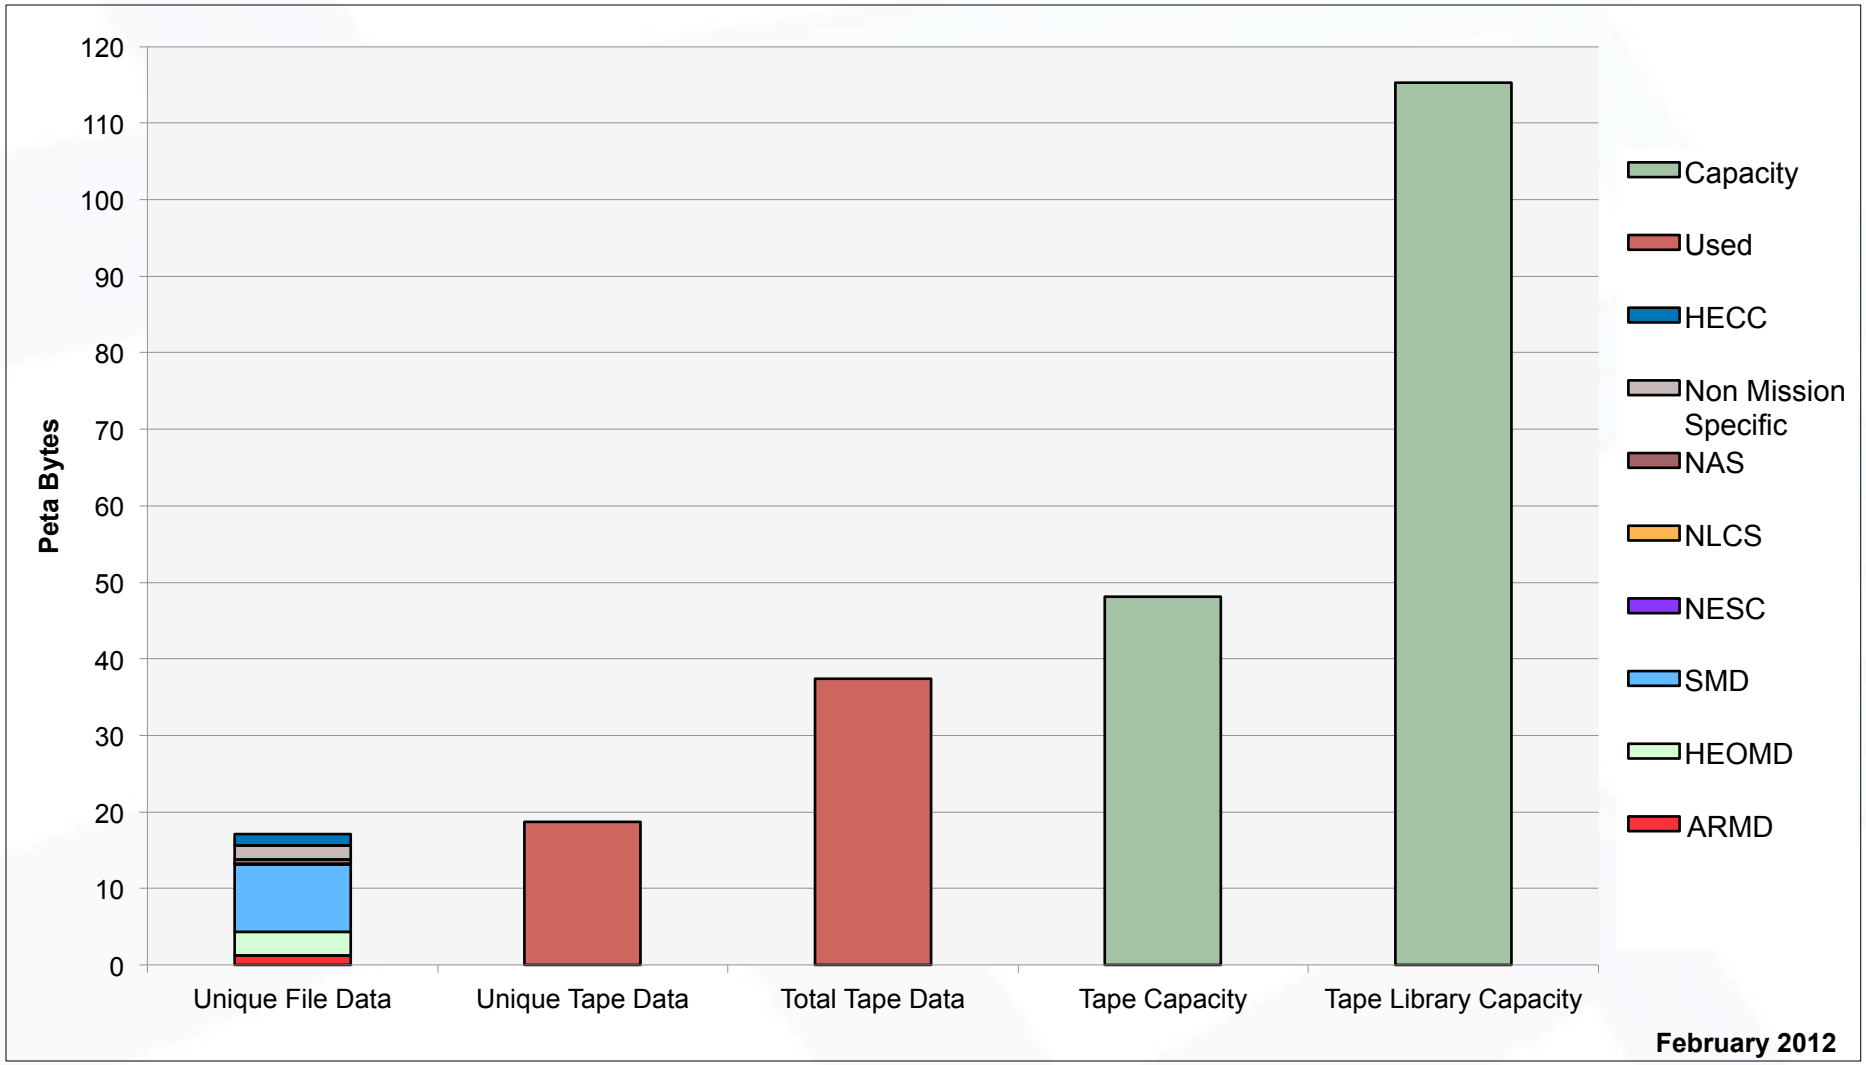

Tape Archive Status

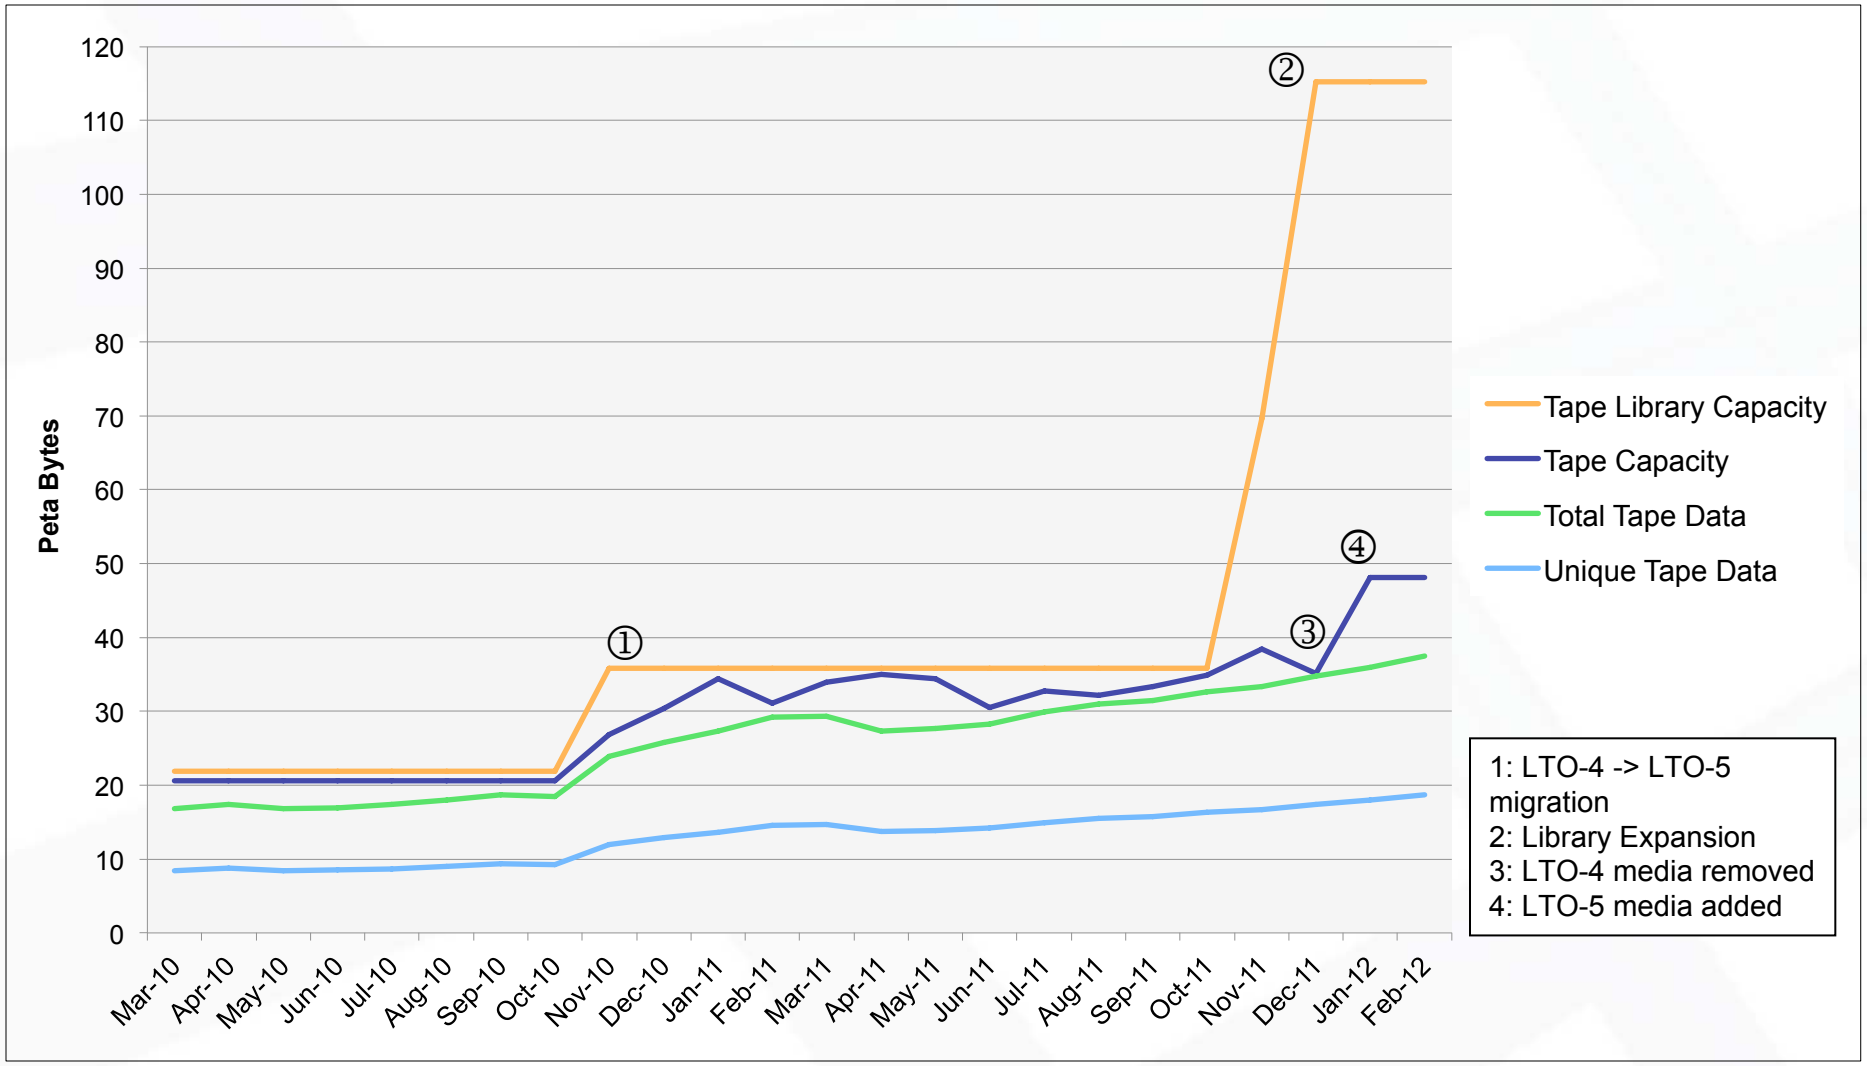

PLEIADES-SPECIFIC CHARTS

Pleiades:
SBUs Reported, Normalized to 30-Day Month

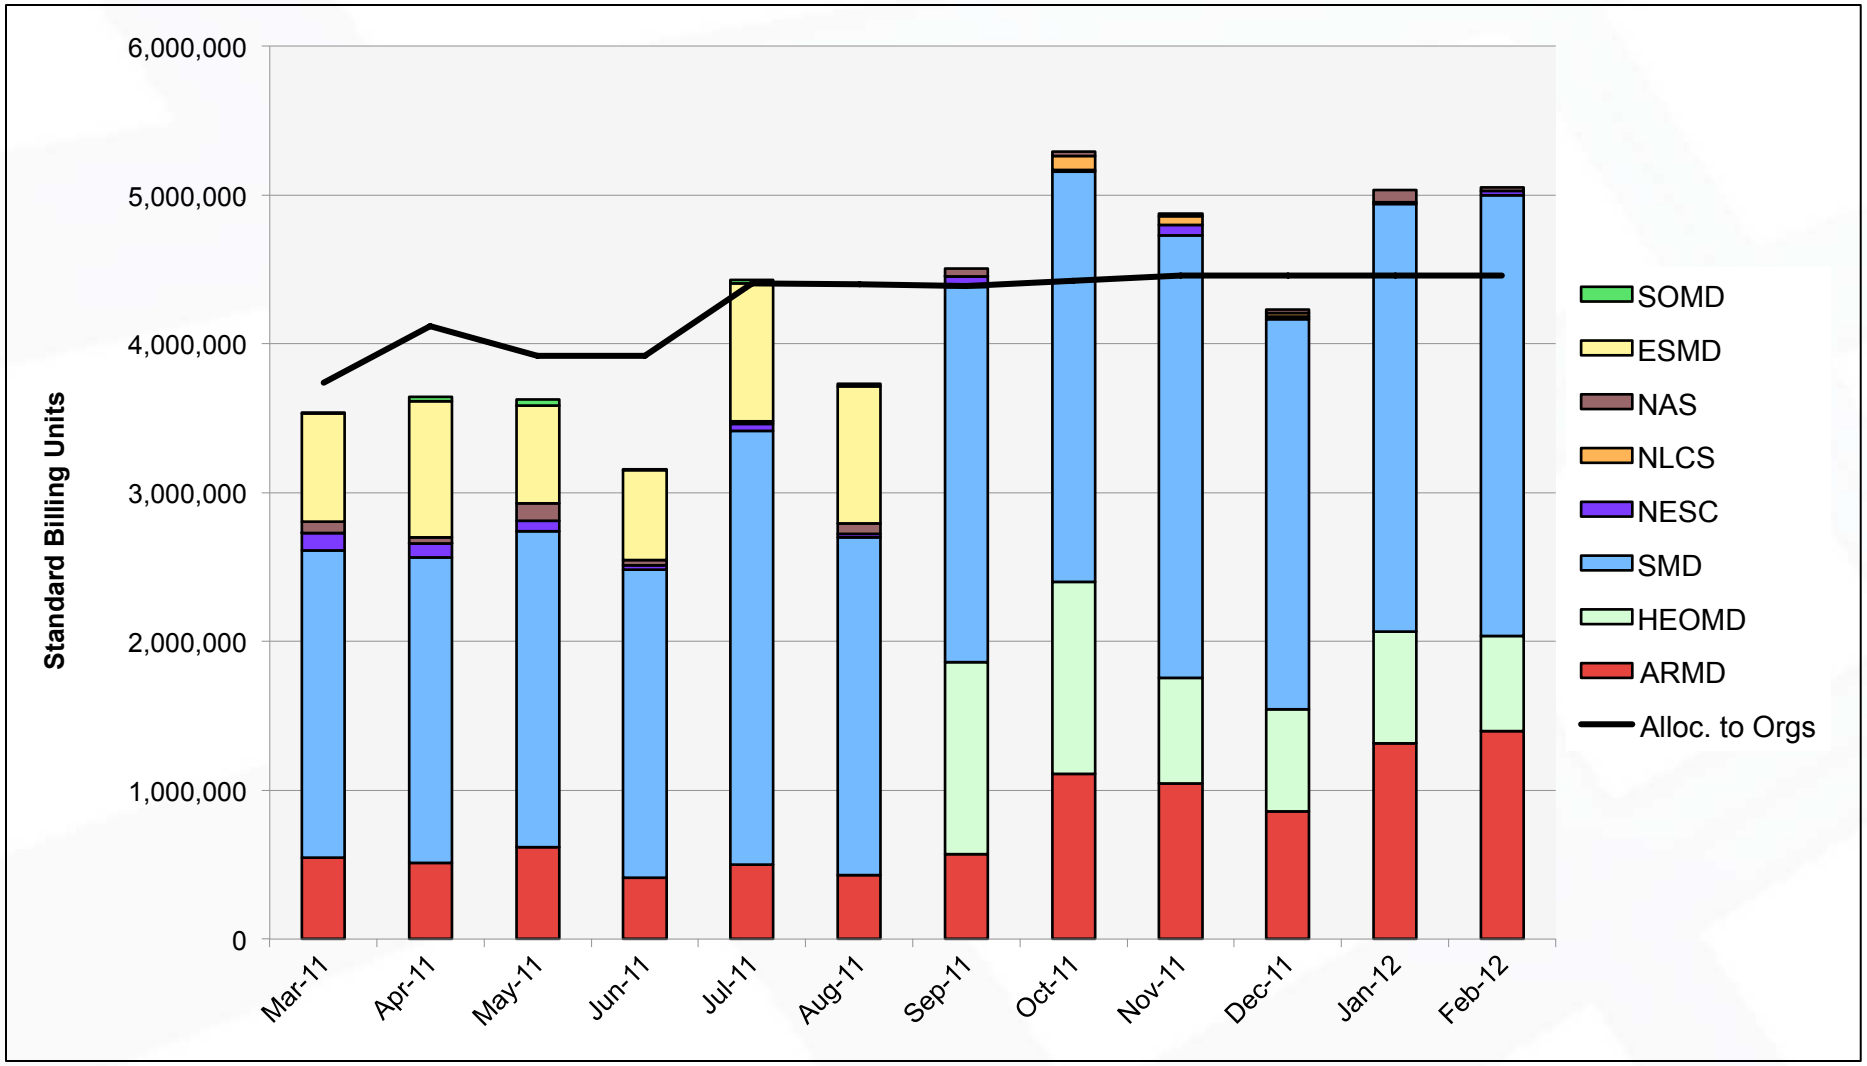

Pleiades:
Devel Queue Utilization

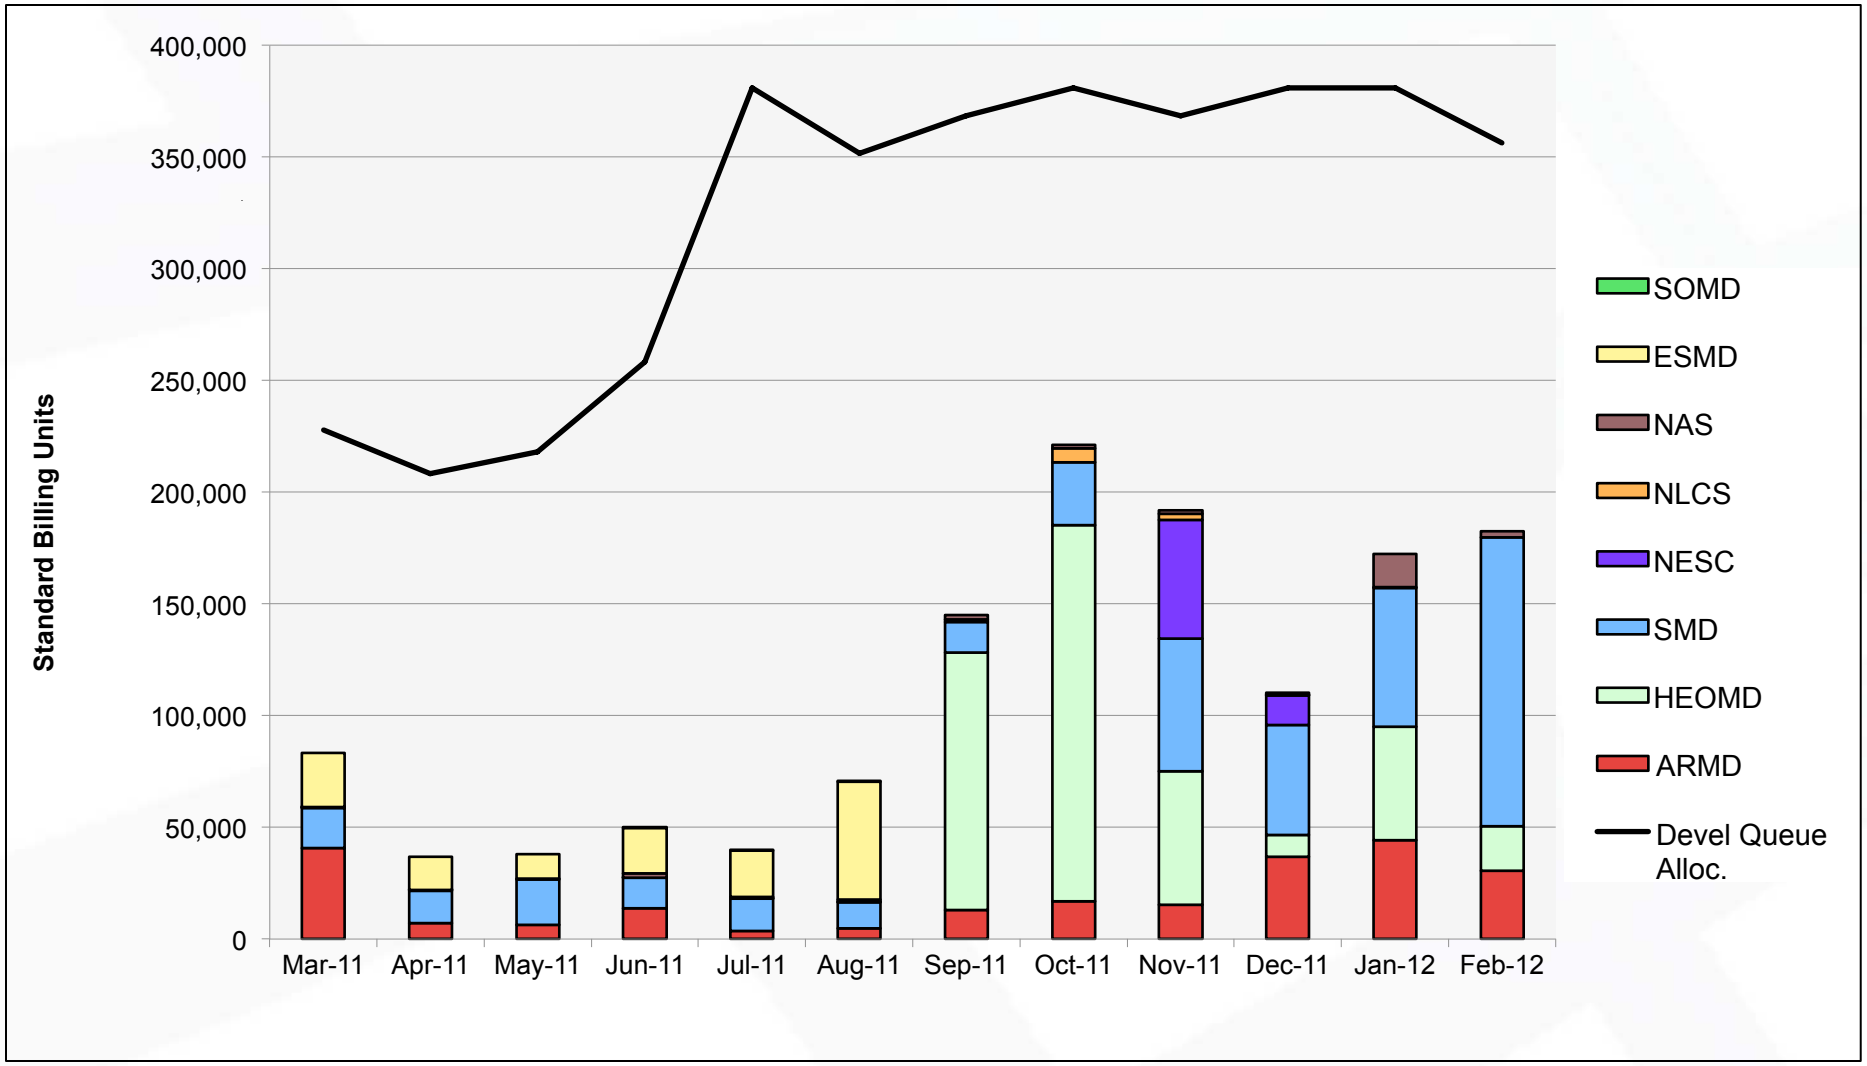

Pleiades:
Monthly SBUs by Run Time

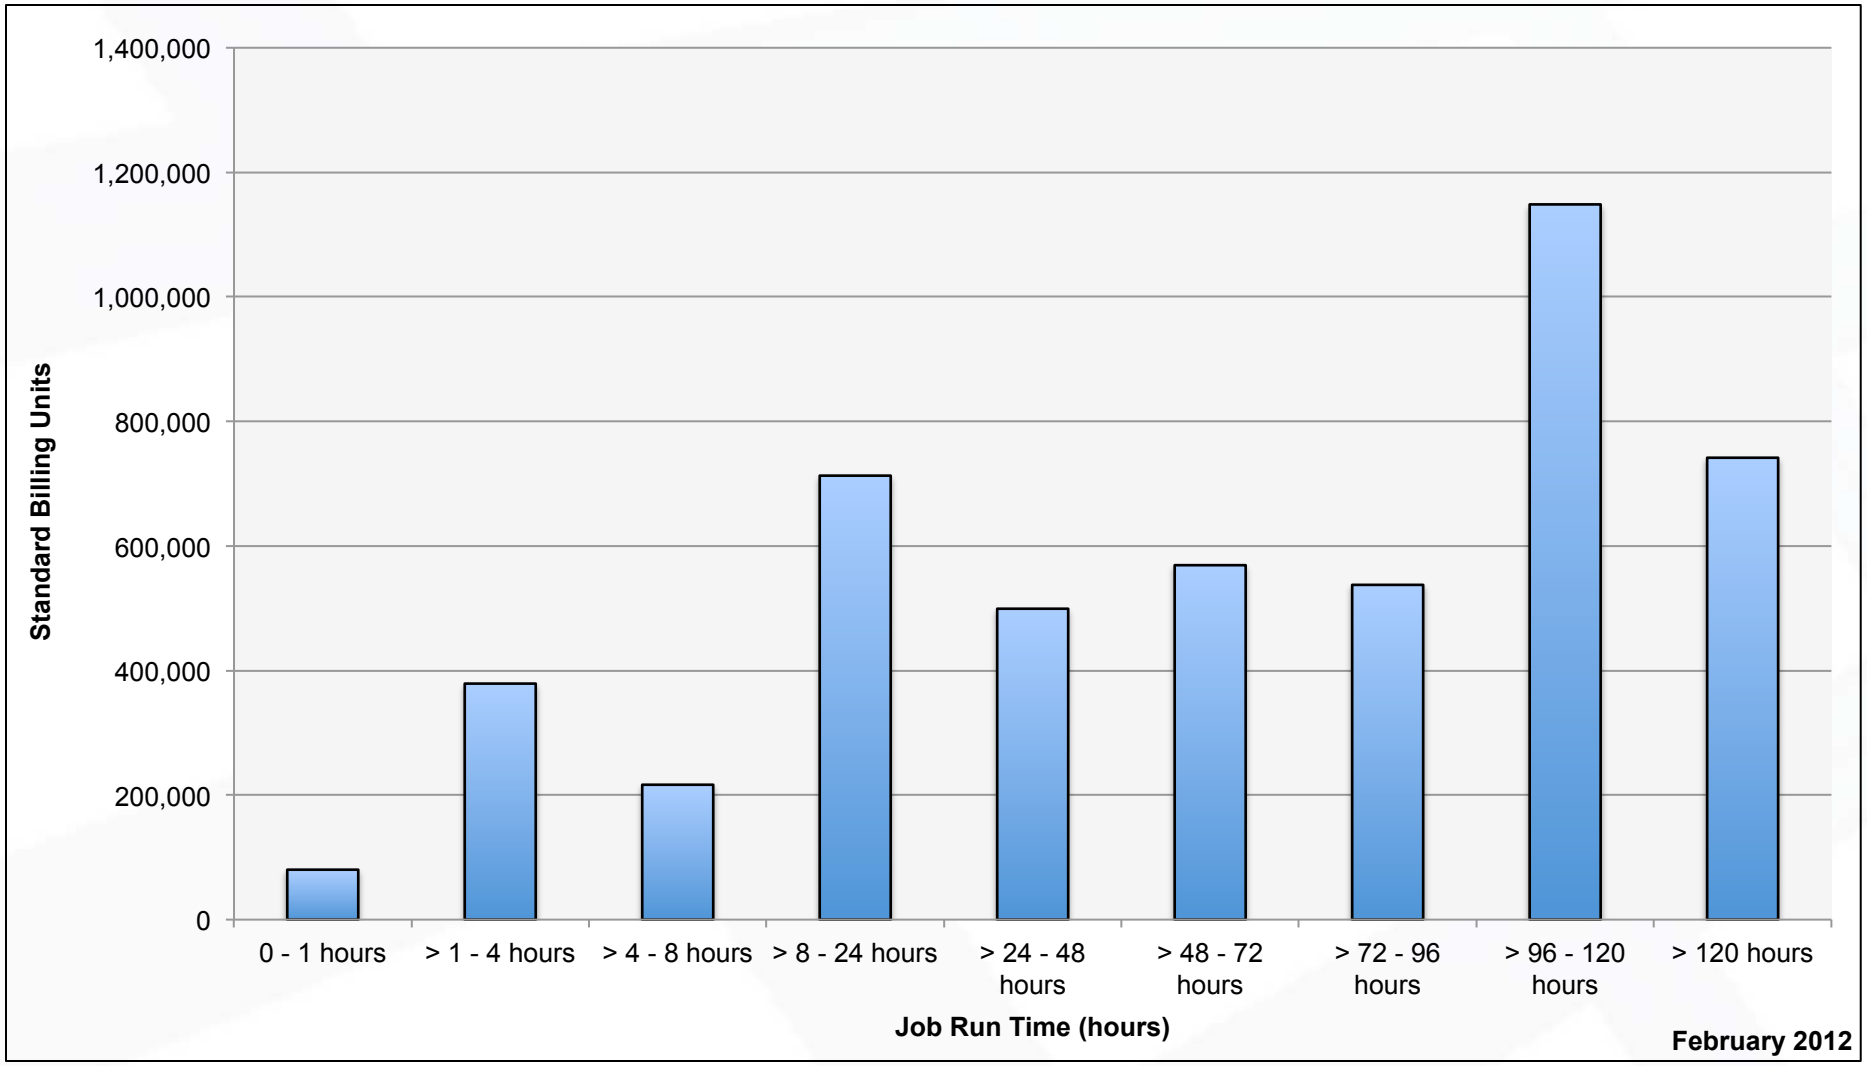

Pleiades:
Monthly Utilization by Size and Mission

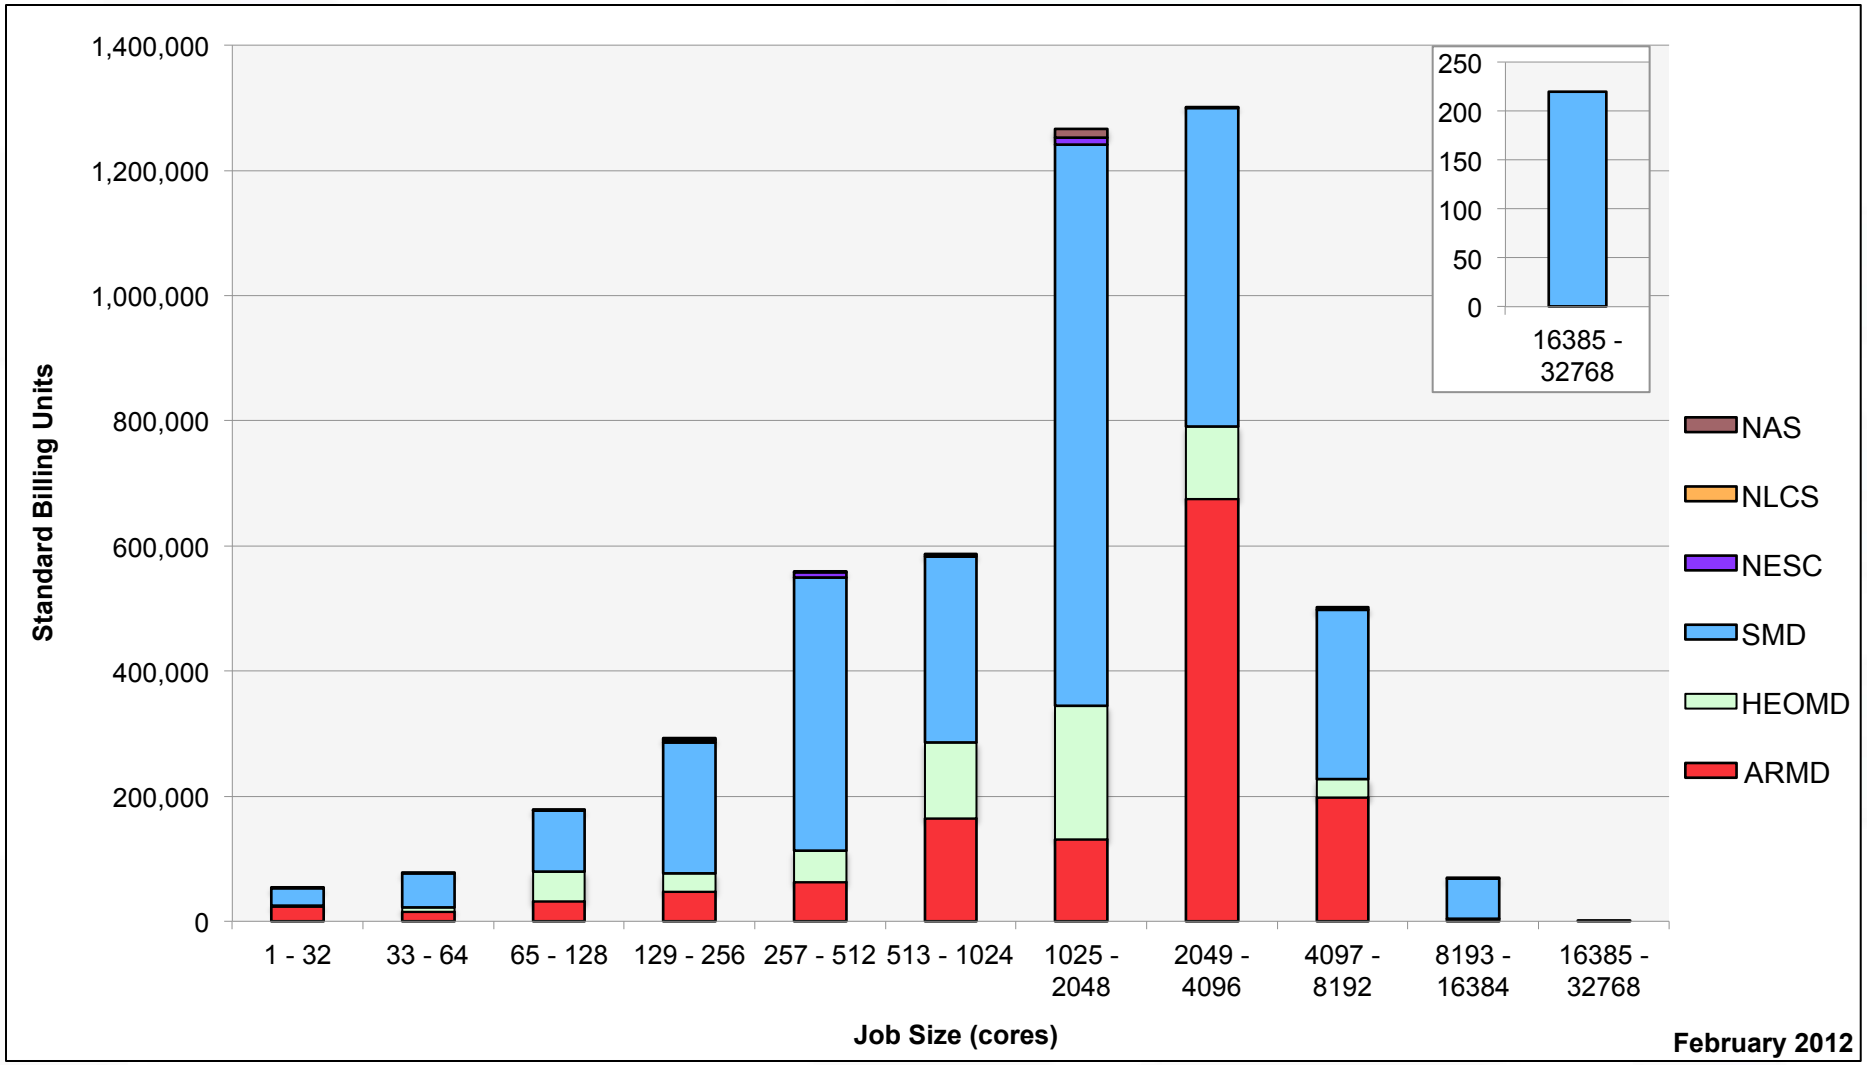

Pleiades:
Monthly Utilization by Size and Length

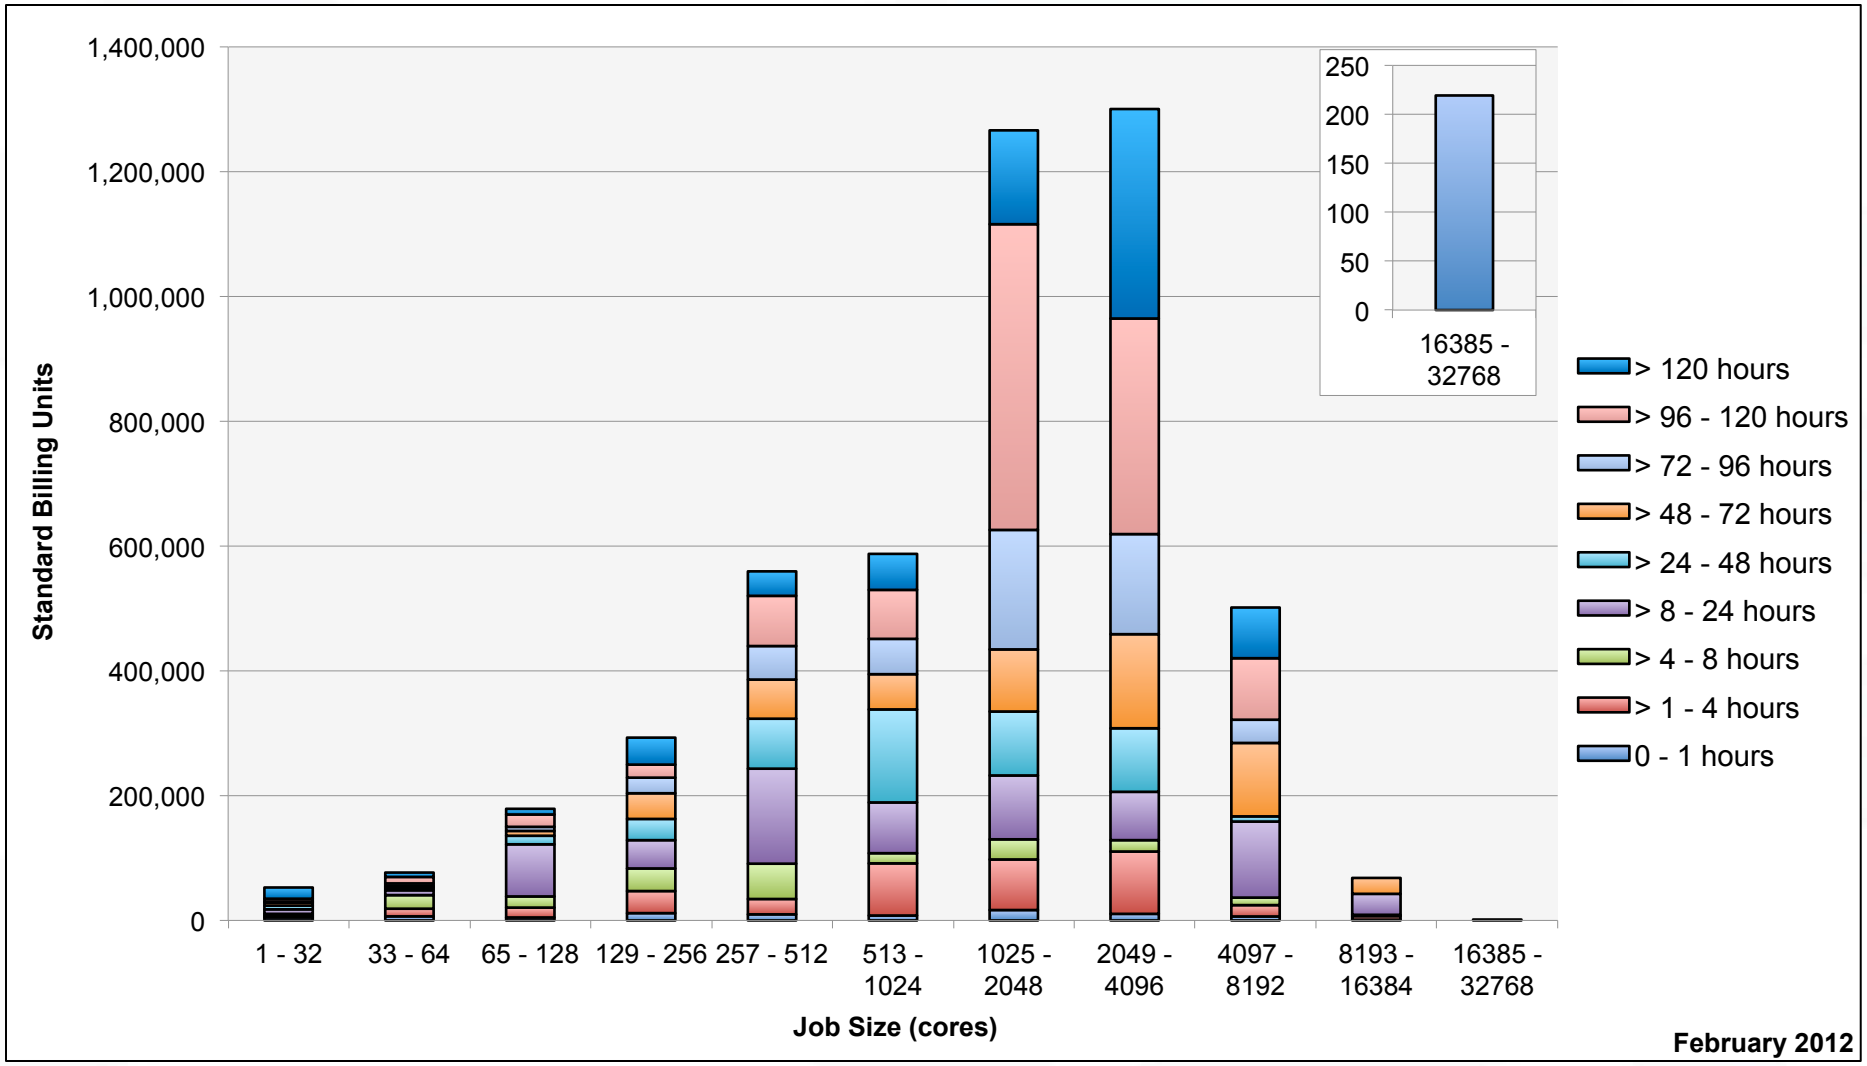

Pleiades:
Average Time to Clear All Jobs

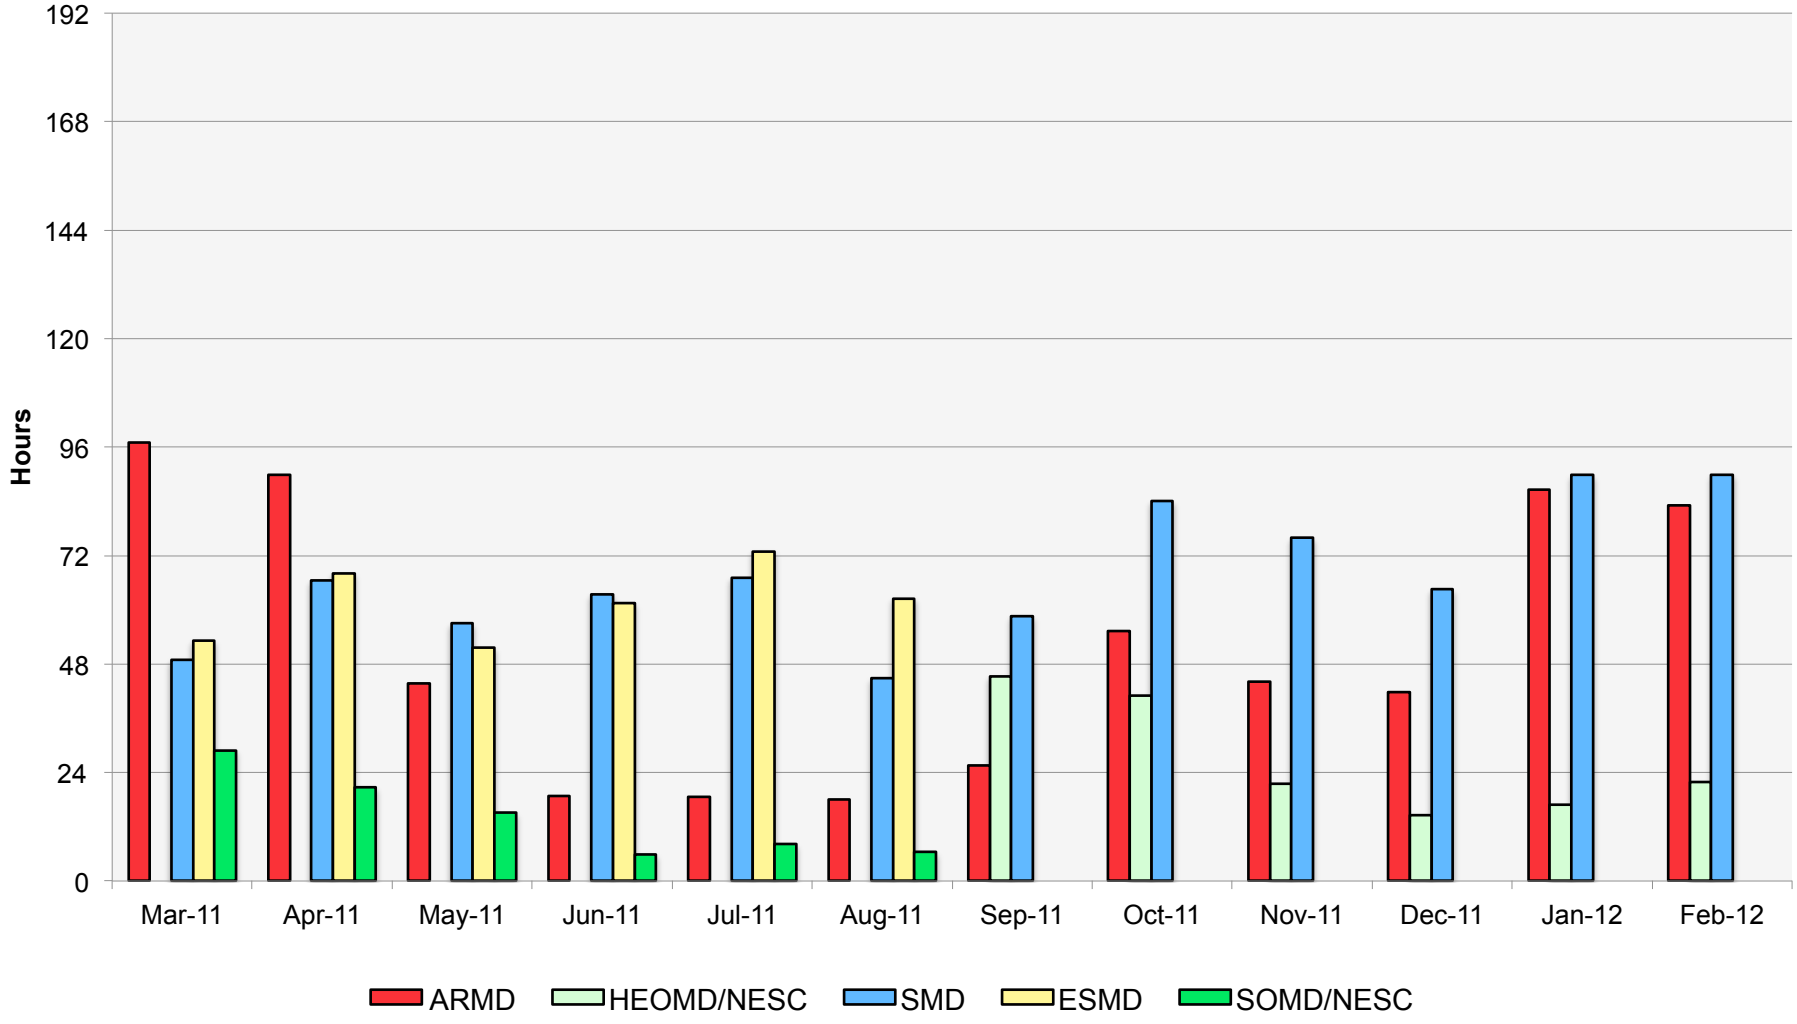

Pleiades:
Average Expansion Factor

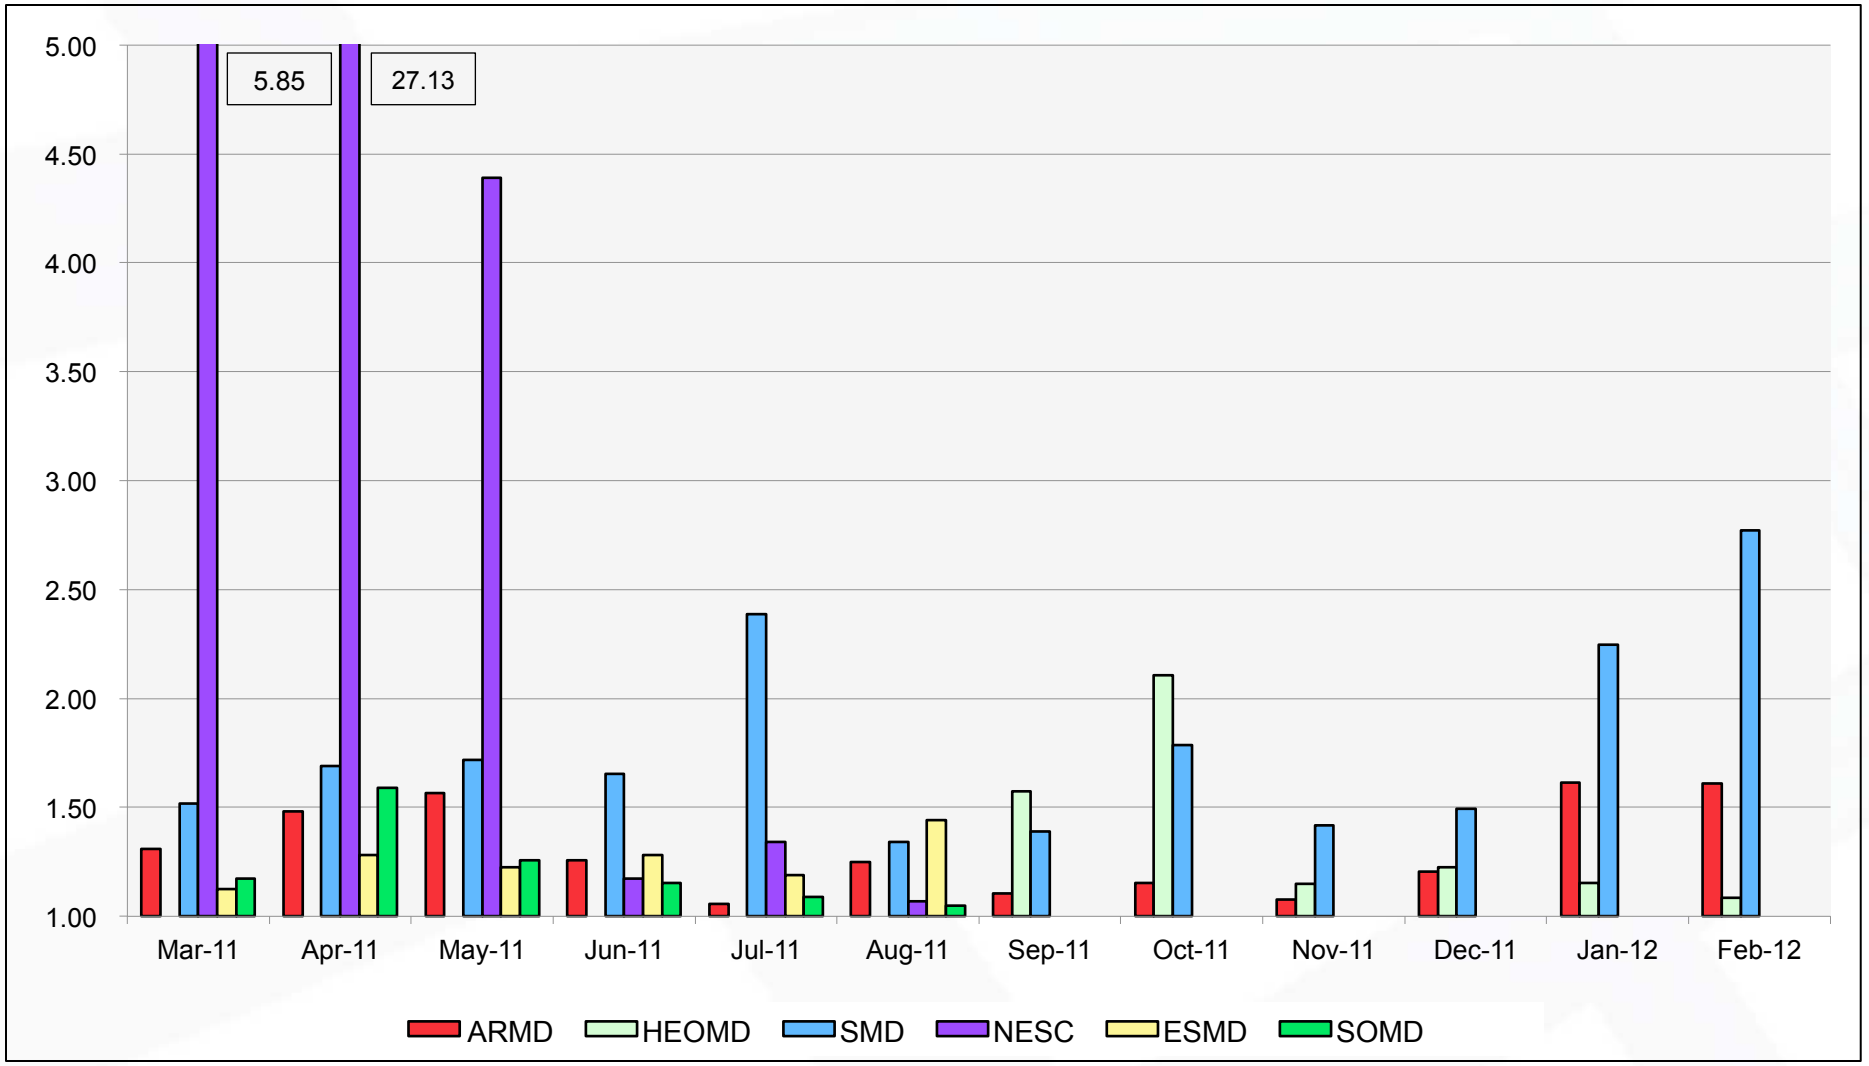

COLUMBIA-SPECIFIC CHARTS

Columbia:
SBUs Reported, Normalized to 30-Day Month

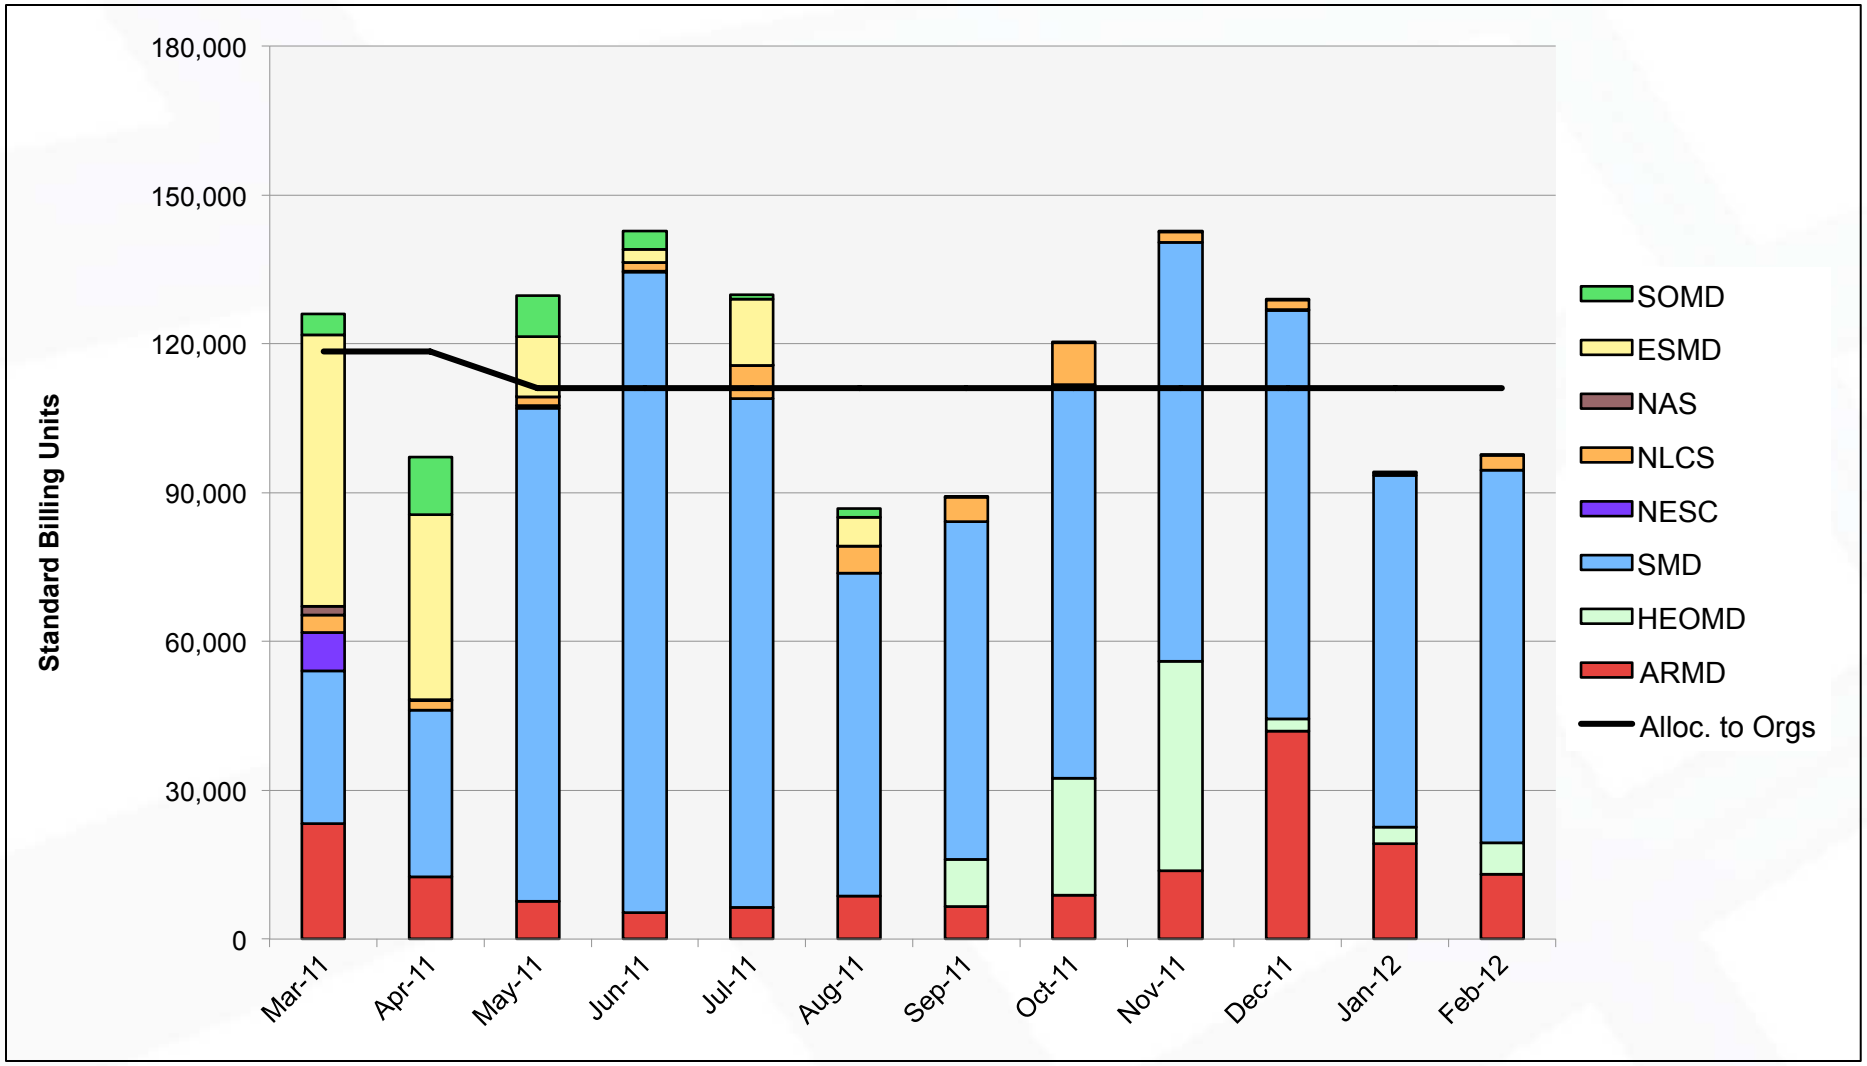

Columbia:
Monthly SBUs by Run Time

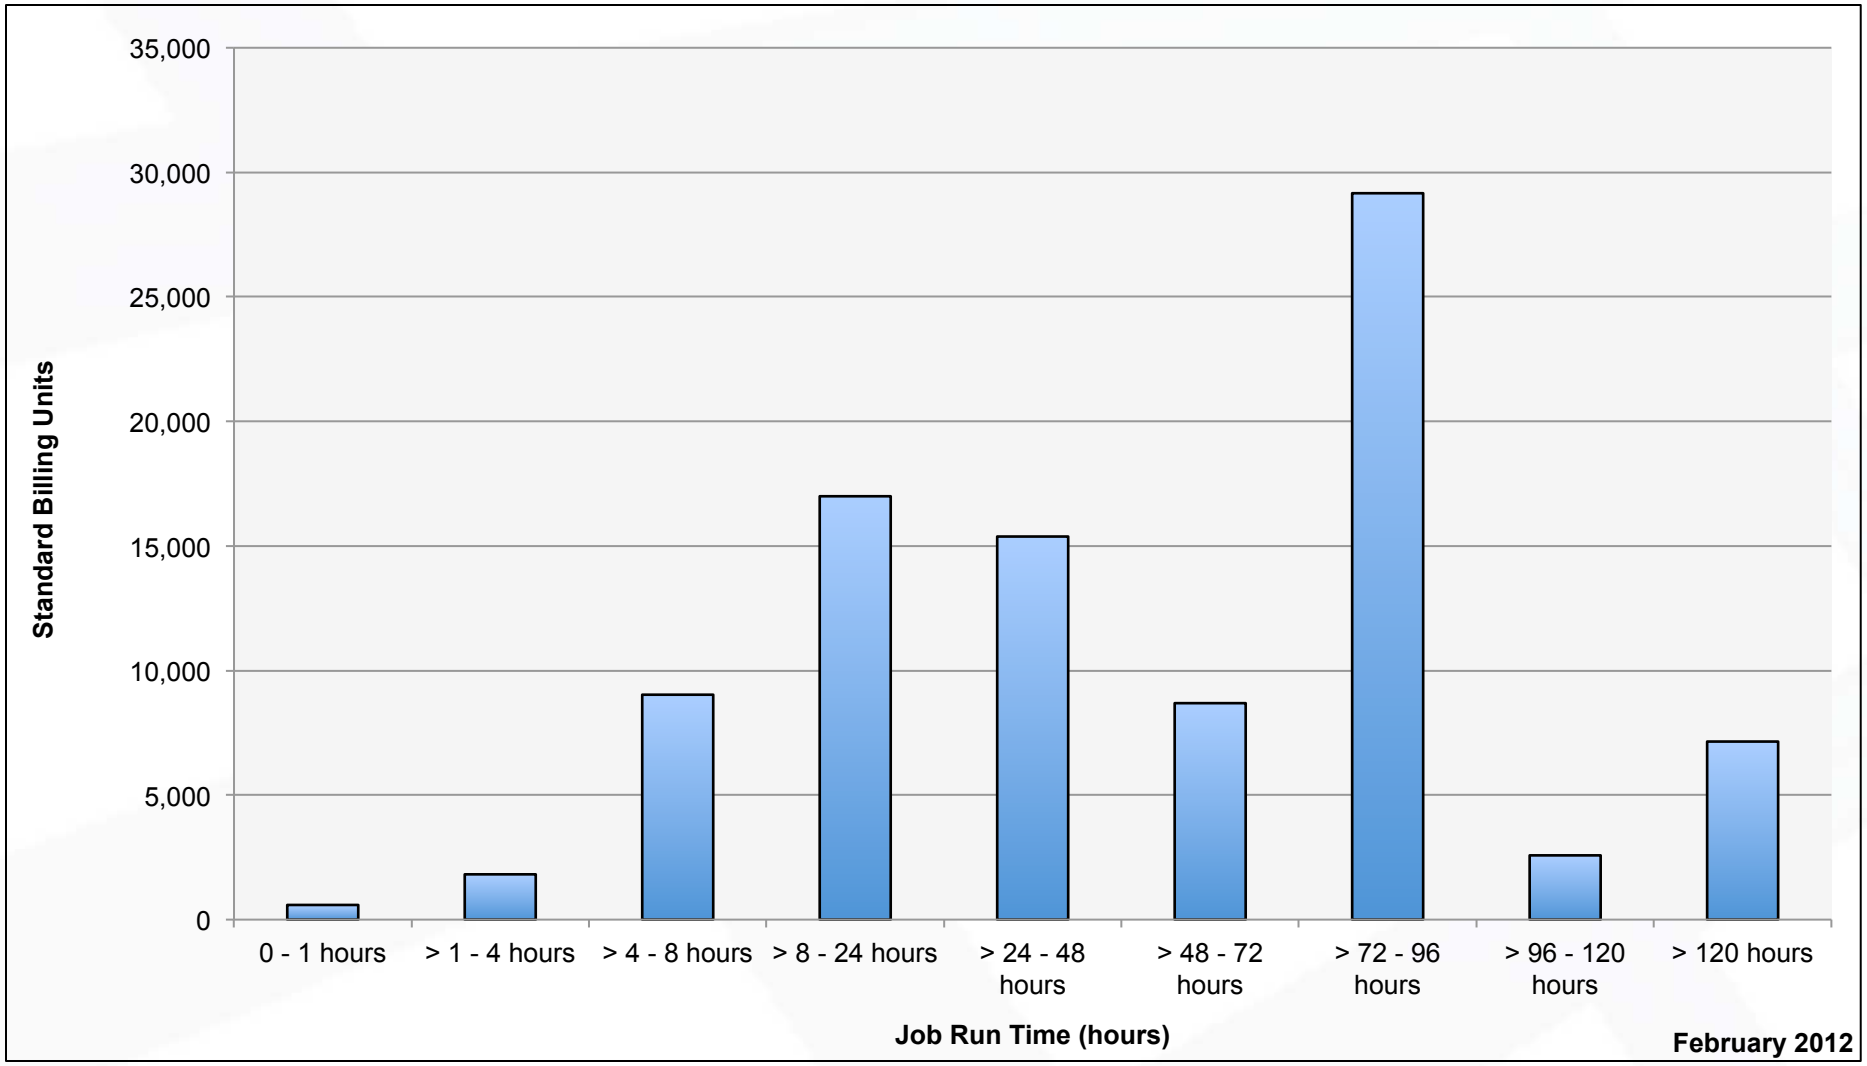

Columbia:
Monthly Utilization by Size and Mission

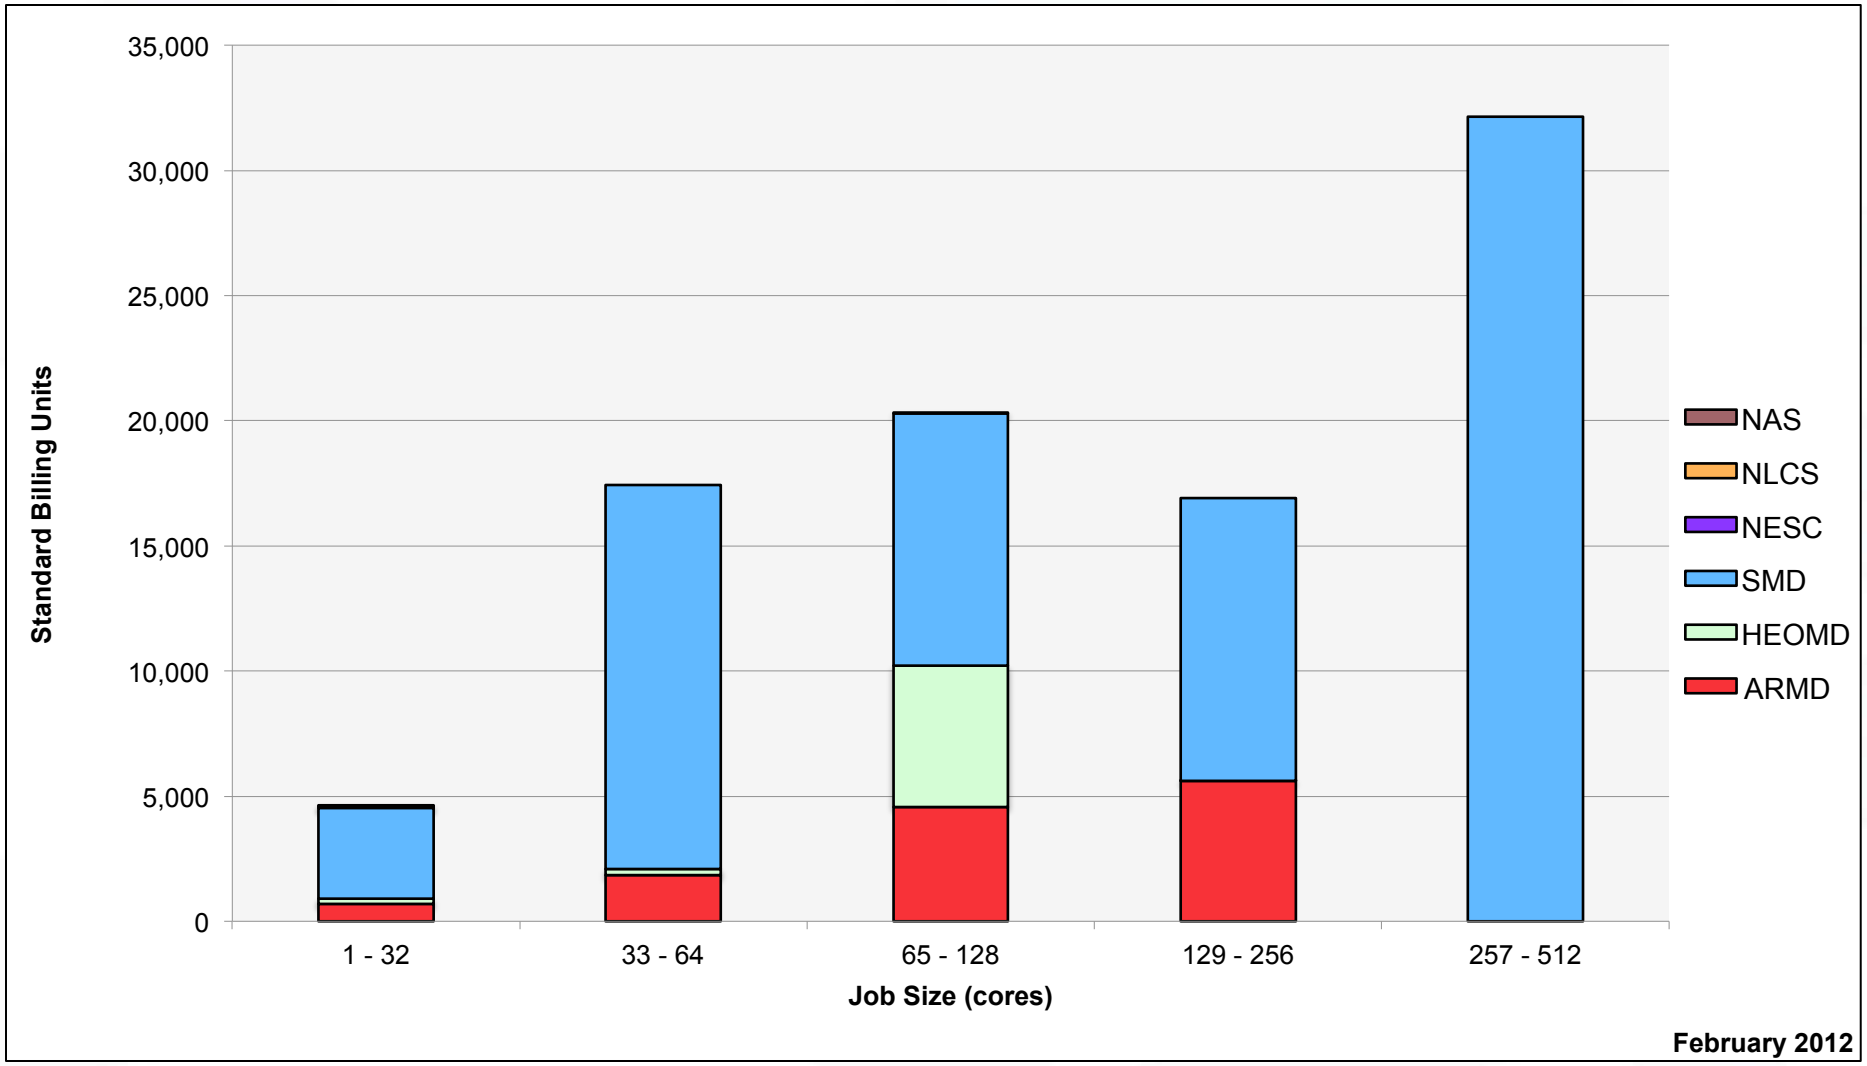

Columbia:
Monthly Utilization by Size and Length

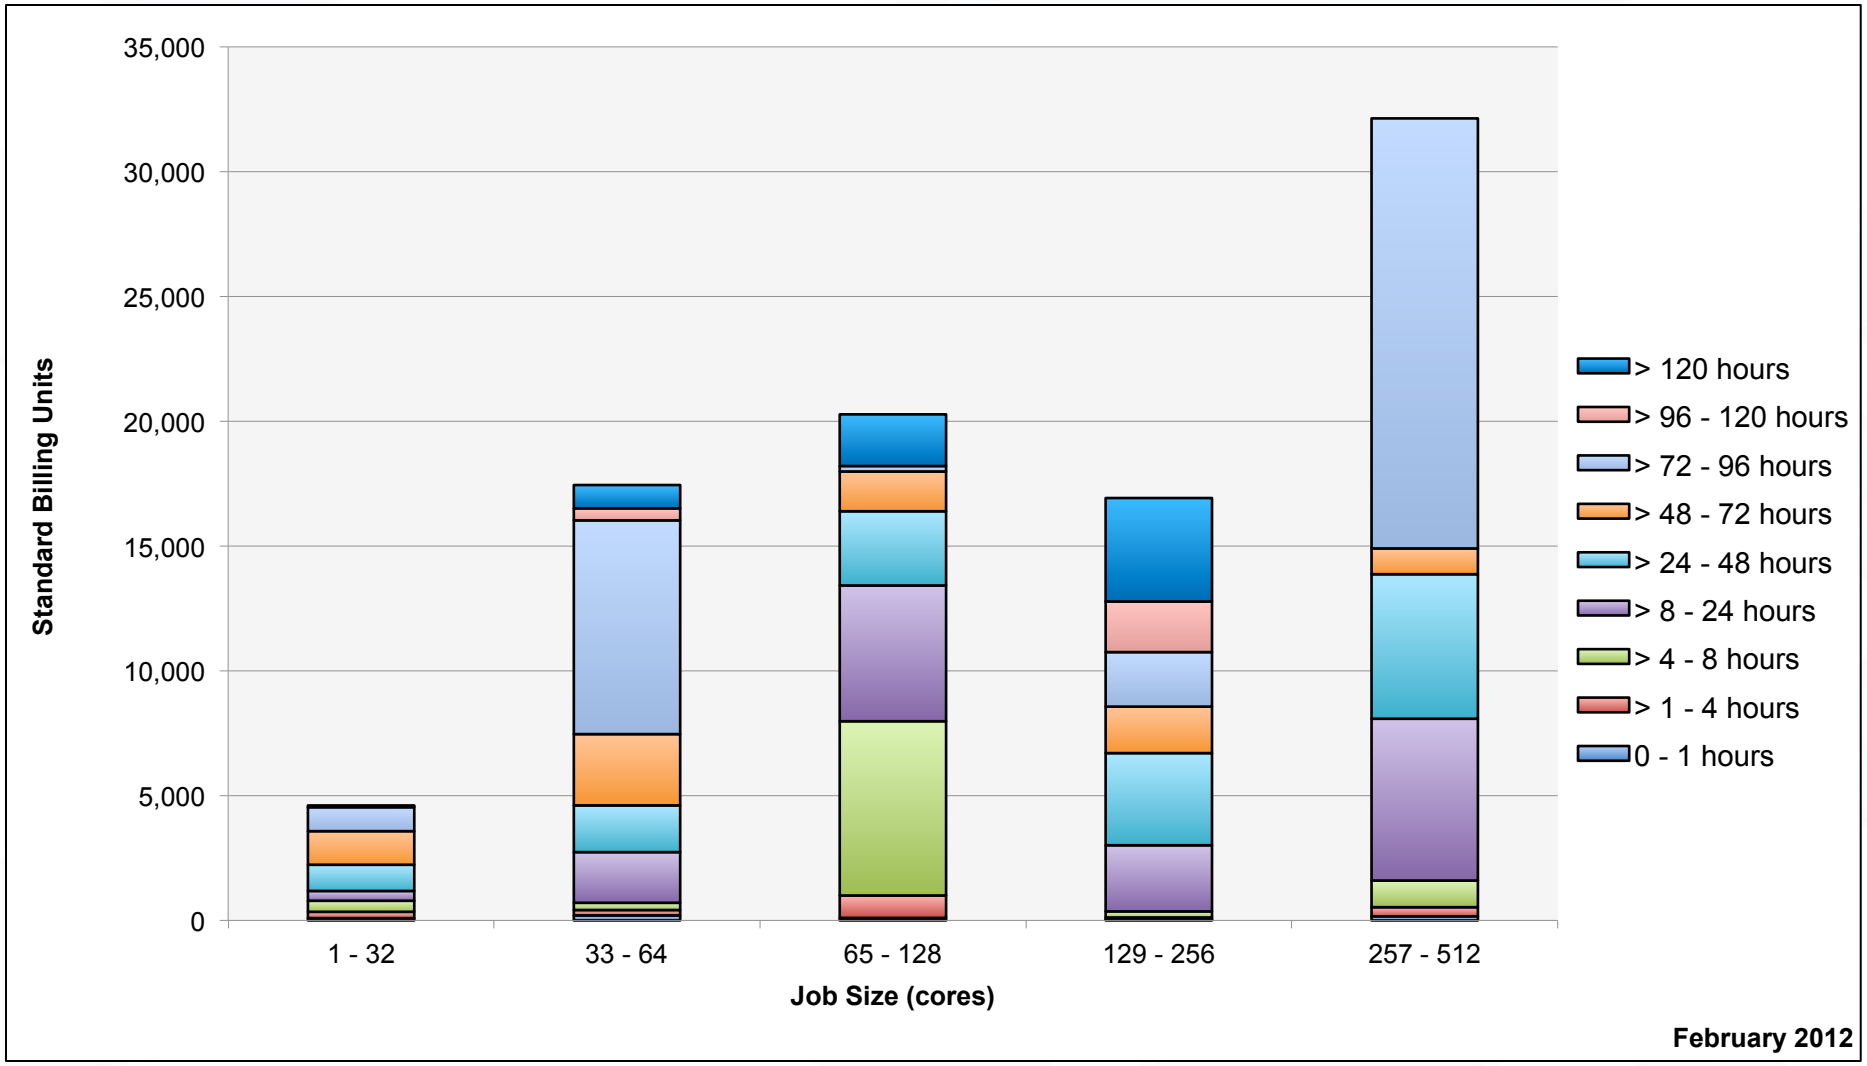

Columbia:
Average Time to Clear All Jobs

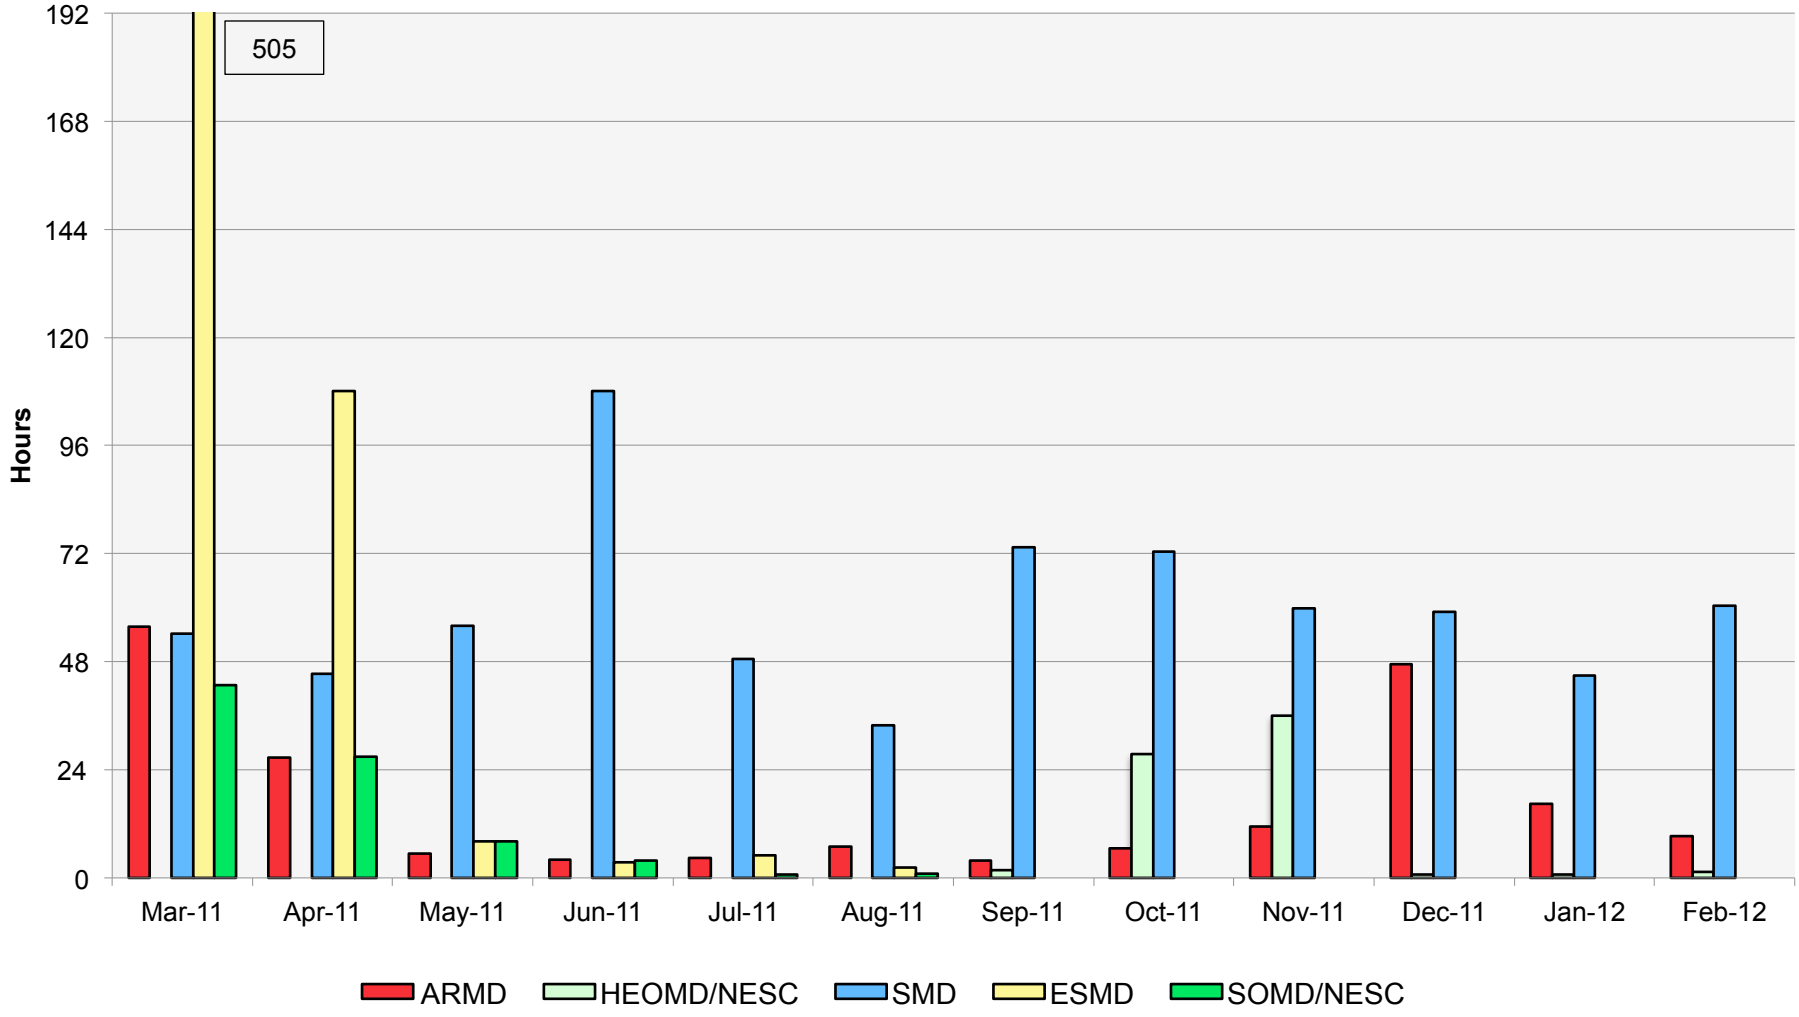

Columbia:
Average Expansion Factor